

The research program of the Center for Economic Studies (CES) produces a wide range of theoretical and empirical economic analyses that serve to improve the statistical programs of the U.S. Bureau of the Census. Many of these analyses take the form of CES research papers. The papers are intended to make the results of CES research available to economists and other interested parties in order to encourage discussion and obtain suggestions for revision before publication. The papers are unofficial and have not undergone the review accorded official Census Bureau publications. The opinions and conclusions expressed in the papers are those of the authors and do not necessarily represent those of the U.S. Bureau of the Census. Republication in whole or part must be cleared with the authors.

**Technology and Jobs:
Secular Changes and Cyclical Dynamics**

By

Timothy Dunne*/***
University of Oklahoma
and U.S. Census Bureau

John Haltiwanger*/**
University of Maryland
and U.S. Census Bureau

Kenneth R. Troske*
Center for Economic Studies,
U.S. Census Bureau

CES 96-7 September 1996

All papers are screened to ensure that they do not disclose confidential information. Persons who wish to obtain a copy of

the paper, submit comments about the paper, or obtain general information about the series should contact Sang V. Nguyen, Editor, Discussion Papers, Economic Planning and Coordination, Center for Economic Studies, Washington Plaza II, Room 211, Bureau of the Census, Washington, DC 20233-6101, (301-457-1882) or INTERNET address snguyen@info.census.gov.

Abstract

In this paper, we exploit plant-level data for U.S. manufacturing for the 1970s and 1980s to explore the connections between changes in technology and the structure of employment and wages. We focus on the nonproduction labor share (measured alternatively by employment and wages) as the variable of interest. Our main findings are summarized as follows: (i) aggregate changes in the nonproduction labor share at annual and longer frequencies are dominated by within plant changes; (ii) the distribution of annual within plant changes exhibits a spike at zero, tremendous heterogeneity and fat left and right tails; (iii) within plant secular changes are concentrated in recessions; and (iv) while observable indicators of changes in technology account for a significant fraction of the secular increase in the average nonproduction labor share, unobservable factors account for most of the secular increase, most of the cyclical variation and most of the cross sectional heterogeneity.

Keywords: technical change, nonproduction labor share, cyclical, secular

This is a revised version of a paper prepared for the April 1996 Carnegie-Rochester Conference on Public Policy.

*We thank Olivier Blanchard, Mark Doms, Robert Fairlie, Peter Klenow, Glenn MacDonald, Bennett McCallum, Allan Meltzer, conference participants at the April 1996 Carnegie-Rochester conference and the May 1996 Society Of Labor Economists conference, and seminar participants at the CES for helpful comments. We also thank Lucia Foster and Catherine Buffington for excellent research assistance.

**Haltiwanger thanks the NSF for support.

***Dunne thanks the Census Fellows program for support.

I. Introduction

Over the course of the past several decades there have been significant changes in the structure of wages and employment in the U.S. economy. Rising wage inequality has been accompanied by an increase in the return to experience over the 1970s and 1980s and an increase in the return to education over the 1980s. Conformable changes in relevant quantities suggest that these price changes reflect changes in the relative demand for skilled workers: employment, participation, and unemployment for workers in the upper half of the wage distribution have been relatively stable while unemployment and withdrawal from the labor force have increased substantially for workers in the lower part of the wage distribution. These changes in the overall U.S. labor market are particularly evident in changes in the structure of the workplace in U.S. manufacturing. Relative to the 1960s, the typical manufacturing worker in the late 1980s is more educated and is more likely to be a professional, manager or technical worker and less likely to be an operator or laborer.¹

¹ Recent research investigating these changes includes Juhn, Murphy and Topel (1991), Davis and Haltiwanger (1991), Katz and Murphy (1992), Juhn, Murphy and Pierce (1993), and Berman, Bound and Griliches (1994). Note that the return to education fell during the 1970s but this has been attributed to the changes in the relative supply of college educated workers over this period (see, e.g., Katz and Murphy, 1992 for extensive discussion of these issues).

The most prominent current explanation for the changing relative demand for skilled workers is that the adoption of new sophisticated capital equipment and the introduction of flexible manufacturing methods has raised the demand for more highly educated and skilled workers. An alternative competing explanation is that there have been product demand changes towards goods requiring more skilled labor (at least in part stemming from an increasingly integrated world economy). While there have been numerous studies in the recent literature seeking to identify the sources of these labor market changes, the verdict is still out since most of the evidence to date is indirect. Much of the existing analysis exploits household data which do not permit direct examination of the alternative explanations of technology adoption or detailed product demand changes. In contrast, in this paper we exploit longitudinal establishment-level data for the U.S. manufacturing sector covering the 1970s and the 1980s which contains a wealth of information about the changing characteristics of individual establishments: employment and wages by worker type (production and nonproduction workers), capital intensity, R&D intensity, detailed indicators of advanced technology adoption, and detailed industry.² These data permit a much more direct examination of the possible explanations for the change in the relative demand for skilled workers.

² Several aspects of our study distinguish our analysis from the relatively few recent studies that exploit the establishment-level data to investigate the connection between technology, wages and employment (see, e.g., Davis and Haltiwanger, 1991; Dunne and Schmitz, 1995; Doms, Dunne and Troske, 1995; and Bernard and Jensen, forthcoming). First, we provide a comprehensive characterization of the timing, heterogeneity and concentration of plant-level changes in the employment and wage structure in the context of investigating the contribution of observable changes in technology. As part of this, a key feature of our study is the documentation and analysis of the strong connection between the low frequency structural changes in the employment and wage structure at the plant with the cyclical dynamics. Second, we investigate the respective contributions of continuing plants and entry and exit. Consideration of the role of entry and exit is important in this setting since many models of technology adoption hypothesize that new technology will be introduced primarily by entering plants that in turn displace exiting outmoded plants.

While our data have a wealth of information about employer characteristics, the primary information on worker characteristics are employment and wages broken out separately for production and nonproduction workers. Using this information, Figure 1 graphs the nonproduction labor share in terms of employment and wage shares and illustrates two key features of the aggregate data. The first is the striking upward trend in the nonproduction labor share. This increase has been interpreted in the recent literature (e.g., Berman, Bound and Griliches, 1994; and Goldin and Katz, 1996) as an important indicator of the overall changes in the structure of the workplace in manufacturing.³ The second key feature of the nonproduction labor share series seen in Figure 1 is that it exhibits an asymmetric pattern over the business cycle. Sharp increases in the share of nonproduction labor during economic downturns are only mildly offset by decreases during recoveries. Consequently, almost all of the long run increase in nonproduction labor share that occurs over the 16 year period, occurs in periods that manufacturing sector employment is contracting. The strong connection between the structural changes and the cyclical patterns in the aggregate data has been neglected in the recent literature but is a fundamental part of our analysis.⁴

Throughout this paper, we follow this recent literature and focus on the nonproduction labor share (both in terms of employment and in

³ The use of industry-level data on production and nonproduction workers to analyze the relative demand for skilled workers has a long history in labor economics. Table 3.7 in Hamermesh (1993) lists over 20 studies using such data for this purpose.

⁴ In this regard, our approach and analysis are in the spirit of the ideas stressed in the recent literature that restructuring and reallocation are concentrated in economic downturns (e.g., Davis and Haltiwanger, 1990; Hall, 1991; Caballero and Hammour, 1994; Mortensen and Pissarides, 1994; and Campbell, 1995. However, in this case, we focus more on the nature of within plant restructuring at business cycle frequencies than on the between plant reallocation of jobs that has been the focus of this recent literature.

terms of wages) as our variable of interest. Although this is dictated by data limitations, we think the use of this variable for this purpose is appropriate for a number of reasons. First, as in the recent literature, we view the nonproduction labor share variable as a basic, though admittedly imperfect, measure of plant-level workforce skill. It is well documented that nonproduction workers are more highly paid and more educated than production workers (Davis and Haltiwanger, 1991; Berndt, Morrison, and Rosenblum, 1992). In addition, the growth in nonproduction labor in manufacturing has been in the higher skill occupations such as professionals (including scientists, engineers, and computer programmers) and managerial occupations, while the decline in production labor has been primarily in lower skill occupations such as operatives and laborers (Davis and Haltiwanger, 1991; Berman, Bound and Griliches, 1994). Second, we view changes in the nonproduction worker share as more broadly reflecting changes in the way plants produce goods. Goldin and Katz (1996) document a variety of changes in the production process over the twentieth century and their effect of the types of workers used in manufacturing production. Similarly, Ichniowski and Shaw (1995) document and analyze changes from assembly-lines to team production and the impact of this change on workplace organization. In addition, Kremer (1993) argues that changes in the complexity of goods produced in turn affect the workforce requirements of the plant. In short, we recognize that plant-level changes in the nonproduction labor share may represent more than simply skill changes in the workforce and we attempt to take this broader perspective into account in interpreting our results.

With these alternative interpretations in mind, this paper examines a number of issues concerning both the microeconomic dynamics of nonproduction labor share changes and the relationship between technology and changes in the skill of the workforce. The first issue

we address empirically is whether the observed aggregate changes in the nonproduction labor share at high and low frequencies reflect a general upward shift in workforce skill (a within plant effect), a reallocation of the employment from continuing low skill to high skill plants (a between plant effect), or entering high skill plants displacing exiting low skill plants (a net entry effect). The results of such basic decompositions can potentially shed considerable light on a variety of competing hypotheses. The hypothesis that trade and other factors have generated a shift in demand towards products that are skill intensive implies that the observed change should primarily be a between plant phenomenon. Demand shifts towards high skill intensive products also potentially have implications for the contribution of net entry. The skill biased technical change hypothesis implies that the observed changes are driven by individual plants retooling their production processes (a within plant phenomenon) or that new technology is introduced by entering plants that displace outmoded exiting plants (a net entry phenomenon). Further, understanding whether new technology is introduced via retooling of existing plants or via the entry of new plants is of fundamental importance in distinguishing between alternative classes of technology adoption models.

The second set of issues we address concerns the timing, heterogeneity and concentration of plant-level changes in the nonproduction labor share. We document the magnitude, concentration, persistence, and cyclicity of the distribution of the plant-level nonproduction share changes. In addition, we document the underlying job creation and destruction dynamics within and between plants by worker type. While we know from recent studies of plant-level employment dynamics that there is tremendous dispersion in the employment growth rate distribution, this need not translate into dispersion in nonproduction share changes. More generally, the

characterization of the distribution is important for understanding the underlying forces driving plant-level changes in nonproduction share. Recent characterizations of plant-level investment (e.g., Doms and Dunne, 1994; Cooper, Haltiwanger and Power, 1995; and Caballero, Engel and Haltiwanger, 1995b) document the importance of lumpy investment spikes in a manner consistent with nonconvexities in capital adjustment costs. For our purposes, nonconvexities in the adjustment costs for adopting new technology (which may be associated with capital adjustment) in the presence of biased technical change in turn imply lumpy adjustment in worker mix at the plant level. In addition, we are interested in the respective contributions of positive and negative changes in nonproduction labor share at the plant level in accounting for the observed aggregate changes. As will become apparent, the observation of very large positive and negative changes in workforce skill, as measured by nonproduction labor share, raises a variety of questions regarding the nature of the bias in technical change.

The last and most important issue we address is the connection between plant-level indicators of technology adoption and changes in the employment structure at both high and low frequencies. We evaluate the contribution of observable factors such as changes in equipment, structures, R&D, and the adoption of advanced manufacturing technologies to high and low frequency changes in the nonproduction labor share. Of particular interest is whether these observable factors can account for the concentration of the long run changes in economic downturns.

The outline of the paper is as follows. The next section considers theoretical issues that help frame the empirical analysis that follows. The intent here is to characterize the relevant driving forces as suggested by the existing theoretical literature. Section III briefly characterizes the data we use. In addition, in this section we characterize the available evidence on the relationship between the

nonproduction labor share measure and workforce skill. Sections IV, V and VI provide a detailed characterization of the high and low frequency plant-level production and nonproduction worker dynamics. Section VII reports the results from our examination of the connection between observable indicators of technology adoption and changes in the nonproduction worker share. Section VIII summarizes our main findings.

II. Theoretical Considerations

A. Within Plant Changes in Workforce Skill

In considering the connection between technical change and workforce skill, we focus on the role of observable indicators of plant-level technology adoption. This focus is motivated by both recent research and the popular perception that the demand for skilled workers has increased relative to that for unskilled workers as manufacturing plants have installed sophisticated capital equipment. While this perspective motivates much of our analysis, we interpret the adoption/retooling decision in a very broad sense. Beyond considering technology adoption via the purchase and installation of sophisticated capital equipment, we view technical change as encompassing a broad range of changes in the production and organizational structure of the plant. The broader interpretation is closely linked to the concept of organizational capital stressed by Hall (1991). Organizational capital reflects the myriad of factors that characterize the production process of an individual plant including the amount and type of capital used, the design and layout of the production process, and the organization of the workforce in teams of managers, office workers, and production workers.

Given this broad definition for changes in "capital," consider the implications for changes in the skill of the workforce at the plant.

Treating skilled and unskilled labor as variable factors of production, the optimal "skill mix" can be determined by short run cost minimization for given output and a given state of technology (where the latter is treated as a quasi-fixed factor for this purpose).⁵ That is, producer *i* minimizes $w_t^s L_{it}^s + w_t^u L_{it}^u$ subject to $y_{it} = F(Z_{it}, L_{it}^s, L_{it}^u)$. Production, y_{it} , is an increasing function of three factors: the "capital" in which the adopted state of technology is embodied, Z_{it} , and two labor inputs (skilled, L_{it}^s , and unskilled labor, L_{it}^u , --treated as variable factors of production). $F(\cdot)$ is assumed to be strictly concave.⁶ The producer takes wages of skilled workers, w_t^s , and unskilled workers, w_t^u , as given.

Optimal skilled and unskilled labor inputs are determined by the standard condition equating the ratio of the marginal products to the ratio of the wages of skilled to unskilled labor, along with the production relationship for given output and Z . For our purposes, it is useful to express the implied optimal skill mix, M_{it} , in a form familiar in empirical analysis:

$$M_{it} = L_{it}^s / (L_{it}^u + L_{it}^s) = m(Z_{it}, w_t^s / w_t^u, y_{it}) \quad (1)$$

⁵ This specification of short run cost minimization is consistent with a fully specified dynamic profit maximization model in which individual producers endogenously adopt new technologies. All we are doing here is characterizing how the variable factors of production are determined for a given state of technology and output. In a fully specified dynamic model determining Z (our index of the state of technology), incurring costs of adoption in the current period would (perhaps with some probability) yield an updated Z in the subsequent period. The costs of adoption may be proportional to current output so that times of adoption affect the scale of operations and in turn the demand for skilled and unskilled labor. This is captured in the short run cost minimization since we are controlling for the level of output. See the discussion below for alternative models of the dynamics of Z .

⁶ The underlying assumption is that there is some additional fixed factor other than Z .

While the skill mix is decreasing in the relative wages of skilled to unskilled workers, the sign and magnitude of the Z-skill complementarity (m_z) depends on the nature of the skill bias in technology adoption. Short run nonhomotheticity (m_y) reflects changes in the skill mix induced by changes in the scale of operations for given Z.

Given non-zero Z-skill complementarity, a key question is, what determines the dynamics of Z. A large recent literature is devoted to endogenizing the technology adoption decision. One important class of models are vintage capital models as in Solow (1960), Chari and Hopenhayn (1991), Cooper and Haltiwanger (1993), Cooley, Greenwood and Yorukglu (1994), and Cooper, Haltiwanger and Power (1995). In these models, Z is characterized as physical capital in which technological progress is embodied. A closely related but distinct class of models characterizes the dynamics of Z via the endogenous innovation and imitation of technologies (e.g., Jovanovic and MacDonald, 1994; Andolfatto and MacDonald, 1993). In all of these models, individual producers must incur costs (both direct and indirect) to acquire and implement new technology. In addition, individual producers are subject to idiosyncratic shocks (e.g., demand, cost, productivity and possibly shocks in the success of adoption). It is the presence of adoption costs, along with idiosyncratic shocks, that implies variation in technology across producers.

A number of factors influence the frequency and timing of adoption at the plant level. First, as noted in the introduction, recent studies of plant-level investment dynamics provide support for the hypothesis that there are nonconvexities in the adjustment costs for capital. In the presence of (potentially related) nonconvexities in the costs of adopting new technologies, technology adoption at the plant level will be lumpy (i.e., infrequent and large changes will be observed). Fixed costs of adoption imply that even with steady improvements in the

leading available technology, individual producers will only infrequently update their technologies. In this spirit, Cooper, Haltiwanger and Power (1995) demonstrate in a model in which plants periodically lock-in at a particular state of technology by installing a specific vintage of capital, the probability of retooling will be increasing in the time since the prior retooling.⁷

An additional source of lumpy technology adoption is erratic improvements in the leading available technology. Major technological breakthroughs in individual industries may be infrequent.⁸ The presence of fixed costs in an environment with erratic improvements in the leading edge technology implies that individual producers may forego minor improvements (or at least delay until improvements have sufficiently accumulated).

A second key factor influencing the timing of adoption is the nature and persistence of the demand and cost shocks that generate fluctuations in profitability for a given state of technology. One reason these shocks are important is that the adoption costs may take the form of lost output or productivity due to the disruption in activity during retooling and reorganization.⁹ The most natural interpretation of these disruption costs is that the plant (or parts of the plant) may need to be shutdown during retooling. Alternatively, in

⁷ This prediction is derived in a setting with an exogenous constant pace of technological progress in the leading available technology. Even in this setting, the adoption cycle will not be deterministic given that plants are subject to idiosyncratic and common shocks. See Proposition 2 in Cooper, Haltiwanger and Power (1995) for a formal derivation of this prediction.

⁸ See Andolfatto and MacDonald (1993) for a discussion of large, infrequent technological improvements. In addition, they incorporate a related source of discrete adjustment by specifying that the probability of successfully innovating is less than one. Thus, individual producers may attempt to innovate for several periods prior to achieving success.

⁹ Formally, this can be modeled as a cost of adoption that is proportional to current output.

interpreting the problem more broadly, these costs could be thought of as the substitution of managerial talent (or other resources) away from production activities to reorganization/retooling activity. The presence of such disruption costs provides a potential rationale for retooling (and associated skill mix changes) to be concentrated in economic downturns since the opportunity cost from the disruption in activity induced by the retooling process is low at such times.¹⁰ Even in the presence of such disruption costs, there still may be incentives for procyclical technology adoption if aggregate shocks to profitability are serially correlated. In the presence of positively serially correlated aggregate shocks, a high current shock to profits will imply higher future profits. This can yield procyclical retooling if the disruption costs are sufficiently small since a producer would prefer to have a new technology available when other factors are generating high profits.¹¹

In short, plant-level skill mix changes will reflect changes in the state of technology, relative wages, and changes in the scale of operations (nonhomotheticity). Changes in the state of technology may be lumpy in the presence of nonconvexities in adjustment costs or erratic improvements in the leading available technology. The timing of lumpy technology adjustment (and associated skill mix changes) will be influenced by a number of factors including the rate of advance of the leading technology, the rate of depreciation of the installed technology, the nature of the costs of adoption and the nature and persistence of shocks.

¹⁰ Arguments along these lines are developed in Hall (1991), Cooper and Haltiwanger (1993), and Cooper, Haltiwanger and Power (1995) in terms of characterizing the response of retooling/reorganization to exogenous aggregate shocks. Andolfatto and MacDonald (1993) make a related argument in a model with endogenous aggregate fluctuations.

¹¹ See Proposition 5 in Cooper, Haltiwanger and Power (1995).

B. Within Plant vs. Between Plant Changes in the Skill Mix

The discussion thus far has emphasized changes in the skill mix within a plant driven by a variety of possible forces. In the aggregate (total economy or industry level), observed changes in the skill mix will reflect within plant changes as well as changes in the employment shares across plants. That is, using the notation from the previous section, changes in the aggregate skill mix can be decomposed as follows:

$$\begin{aligned} \Delta M_t = & \sum_{\text{continuers}} (L_{it-1}/L_{t-1})\Delta M_{it} + \sum_{\text{continuers}} (M_{it-1}-M_{t-1})\Delta(L_{it}/L_t) + \sum_{\text{continuers}} \Delta(L_{it}/L_t)\Delta M_{it} \\ & + \sum_{\text{entering plants}} (L_{it}/L_t)(M_{it}-M_{t-1}) - \sum_{\text{exiting plants}} (L_{it-1}/L_{t-1})(M_{it-1}-M_{t-1}) \end{aligned} \quad (2)$$

where (consistent with the notation used above) M_t represents the aggregate skill mix in period t (in practice, we use the ratio of nonproduction worker employment to total employment in the exercises which follow), L_t represents aggregate total employment and the corresponding plant-specific terms are defined accordingly. The first term represents the within plant component for continuing plants between period $t-1$ and t , the second the between plant component for continuing plants, the third term a covariance component for continuing plants and the last two terms reflect the contribution of entry and exit, respectively.¹² Much of our discussion thus far has referred to

¹² This decomposition is closely related to, but distinctly different from, the decomposition used by Berman, Bound and Griliches (1994) among others. Their decomposition involves no covariance term since they used average (across time) employment shares in the within component and average skill intensities in the between component. We have chosen to represent the decomposition in this alternative fashion for two reasons. First, our decomposition easily incorporates the role of entering and exiting plants while the alternative does not. Second, we think it is interesting to consider separately the contribution of the covariance component. Note that in our decomposition, for the between and the net entry terms, each component is deviated from the overall initial average skill mix. Thus, the increase in the employment share of a plant contributes positively to the overall change

(continued...)

the first component of this decomposition: the within plant component of changes in the skill mix.

The between plant component arises from the job reallocation across plants induced by sectoral and idiosyncratic cost, demand, or productivity shocks. The above decomposition makes clear that between plant job reallocation is only relevant for changes in the aggregate skill mix if the underlying forces changing total plant-level employment are correlated with the skill mix at the plant. Thus, for example, if the demand for products produced by technologies that are skill intensive increase disproportionately then this can lead to an increase in the aggregate skill mix even if there are no changes in within plant skill intensities. This may have arisen from increased international competition that shifted U.S. production away from low skill intensive products to high skill intensive products.

The process of technology adoption itself is a factor linking between and within plant changes in the skill mix. For example, if adoption is skill biased and adoption leads to an increase in employment, then these combined effects will produce a positive covariance. More generally, however, the adoption of technology will have industry and general equilibrium effects that generate both between and within plant skill mix changes. First, consider the implications for changes in relative wages induced by biased technical change. Suppose for the moment that most industries exhibit skill biased technical change. This skill biased technical change will increase the demand for skilled labor and potentially increase the relative wages for skilled workers. For plants in industries without skill biased technical change or for plants that have not adopted the latest technology, the change in the relative wages will induce a

¹²(...continued)
only to the extent that it has higher than the average initial skill mix.

decrease in the skill mix. For plants that are adopting new technology that is skill biased, the relative wage change will dampen their skill mix change.

Another relevant industry equilibrium effect to consider is that with stable industry demand, technological progress will yield a falling industry price. Depending on the elasticity of industry demand, this may yield a decrease in industry employment. In terms of the within industry dynamics, the impact is analogous to the between plant effects discussed above (e.g., plants that have not adopted will decrease their employment share).

In short, there will be an endogenous evolution of the cross sectional distribution of technologies (or equivalently, a cross sectional distribution of vintages of "capital") within the same industry in this class of models. The presence of idiosyncratic shocks and adoption costs implies that not all plants will have adopted the latest technology in any given period. Further, common aggregate shocks as well as the distribution of idiosyncratic shocks affect the evolution of the cross sectional distribution.¹³ The evolution of the cross sectional distribution of technologies in the presence of biased technical change in turn generates rich dynamics in the cross sectional distribution of skill mix changes (the within plant component) and employment shares (the between plant component) which we attempt to characterize empirically.

C. Entry and Exit

Another potentially important contributing factor that appears in the above decomposition is the role of entry and exit.

¹³ Andolfatto and MacDonald (1993) emphasize learning externalities that may act to induce bunching of technology adoption as another factor influencing the evolution of the cross sectional distribution.

The contribution of net entry may reflect a variety of alternative factors. First, a large class of relevant models (e.g., Campbell, 1995; Caballero and Hammour, 1994; Lambson, 1991) points towards entry as being the primary way in which new technology is introduced into the economy. This class of models is similar in spirit to the within plant technology adoption models discussed above. In this type of model, new plants incur a fixed cost to adopt the latest technology and in turn old plants with outdated technologies are induced to exit. If technology is skill biased, then skill mix changes will be observed via the entry and exit process.

A second factor influencing the contribution of the net entry component of equation (2) is that changes in product demand will imply differential patterns of net entry across industries. If product demand changes are correlated with the skill intensities of the production processes (e.g., demand for high skill products increases), then this will yield a systematic contribution of net entry to the aggregate skill mix change. In our empirical analysis, we characterize the respective contributions of entry and exit to changes in the skill mix and attempt to distinguish between the various interpretations of the contribution of entry and exit.

D. Other Frictions

The models discussed thus far emphasize one type of friction, retooling costs associated with introducing a new technology or fixed costs of opening a new plant. Since we are interested in exploring the implications of technological adoption for job dynamics, it is important to emphasize other frictions that may be relevant for these changes. The above

discussion treats both skilled and unskilled labor as variable factors of production. This assumption seems reasonable given that our data are at an annual frequency but even at an annual frequency some labor adjustment costs (e.g., search, hiring, firing) may still be relevant.¹⁴ This is especially important to the extent that these frictions differ by the skill type of workers. Since it is often presumed that adjustment costs are higher for skilled workers, any high frequency changes that we observe in the skill mix may reflect these differential employment adjustment costs rather than the factors we have emphasized. In terms of our analysis, these adjustment costs act as a form of nonhomotheticity in the high frequency fluctuations in the skill mix.¹⁵ In our empirical work, we attempt to distinguish between these alternative explanations of changes in the skill mix at high frequencies.

III. Data Description, Measurement Issues and An Empirical Road Map

The data used in this study come from the Longitudinal Research Database (LRD) which is a compilation of the plant-level data from the Census of Manufactures (CM) and the Annual Survey of Manufactures (ASM) for the period 1972 to 1988. For each plant-year observation, the data contain detailed information on production and cost variables such as employment, shipments, and capital investment.

A. Nonproduction Labor Share as a Measure of Skill

¹⁴ Although the empirical work on labor adjustment costs suggests they are most relevant for monthly and quarterly data (e.g., Hamermesh, 1993; and Caballero, Engel and Haltiwanger, 1995a).

¹⁵ A closely related form of nonhomotheticity could arise if there is overhead labor that consists mainly of nonproduction workers.

As we briefly discussed in the introduction, a key issue in this paper is how we measure workforce skill. The data allow us to disaggregate employment into two types of workers -- production workers and nonproduction workers. Production workers include workers "engaged in fabricating, processing, assembling, inspecting, receiving, storing, handling, packaging, warehousing, shipping, maintenance, repair, record keeping, janitorial, and guard services up through the line-supervisor level." Nonproduction employees are all other workers including "sales, sales delivery, clerical, management, professional, technical employees and construction employees." (U.S. Census Bureau (1991), p. A-1). Given these worker classifications, we construct two different plant-level measures of skill. The first is the ratio of nonproduction workers to total employment. The second, is the ratio of wages paid to nonproduction workers to total payroll. This latter measure is of interest for a number of reasons. First, the changes in the demand for more skilled labor may be reflected in changes among production and among nonproduction workers rather than in shifts between production and nonproduction workers. Examining the nonproduction worker wage share provides some perspective on this problem. Second, a shift in demand towards skilled workers will be understated in examining the ratio of nonproduction workers to total employment given that the implied increase in the skill premium will dampen the employment changes. Third, the cost share emerges as the relevant dependent variable when considering specific functional forms (e.g., translog) of the cost function. Since we exploit such specifications in the analysis in section VII it is useful to characterize the skill mix changes from this perspective as well.

One obvious question is whether these two alternative measures (employment and cost share based) of the nonproduction labor share are reasonable measures of skill. One piece of

evidence that suggests the production/nonproduction worker distinction is closely linked to skills is the wages paid to these workers. To the extent that the labor market in U.S. manufacturing can be viewed as competitive, then workers with a higher marginal product (more skill) should receive higher wages. Davis and Haltiwanger (1991), using data from the LRD from 1963 to 1986, show that the average wage of nonproduction workers is \$12.86 per hour, while the average wage of production workers is \$8.56 per hour. Davis and Haltiwanger also show that the difference between nonproduction worker and production worker wages rises by 29 percent between 1975 and 1986. These numbers suggest that, at least based on wages, nonproduction workers are more skilled than production workers, and also the return to being a nonproduction worker has increased over this period (which may reflect either an increased skill premium or increases in the relative skills of nonproduction workers).¹⁶

An alternative, commonly used, measure of skill is education. To see whether plants with relatively more nonproduction worker labor also employ more educated workers, we use data from the Worker-Establishment Characteristics Database (WECD) to examine the educational distribution of workers within plants. The WECD is a cross-sectional employee-employer matched database created at the U.S. Census Bureau (see Troske (1995) for details). Using these data we find that only 7.4 percent of workers in plants in the lowest nonproduction share quartile have

¹⁶ One factor that may be important in considering these wage differences is if workers use different means to acquire human capital. For example, if nonproduction workers acquire most of their human capital by attending school, while production workers purchase training from their firm through lower wages, then these wage differentials may simply reflect the difference in how workers pay for human capital. However, evidence from the training literature suggests that this is not the case. First, white collar workers are more likely to receive training than blue collar workers (Amirault, 1995; Altonji and Spletzer, 1991). Second, even among production workers, it is the most educated production workers who are the most likely to receive training (Lynch and Black, 1995).

Bachelors degrees while 21.1 percent of workers in plants in the highest nonproduction share quartile have Bachelors degrees. In addition, Berman, Bound and Griliches (1994) using data from the 1987 CPS show that only 17 percent of blue collar workers have more than a high school degree, compared with 35 percent of clerical workers, 70 percent of sales workers and 78 percent of managers and professionals. Finally, Davis and Haltiwanger (1991), using data from the CPS from 1973 to 1987 show that, in U.S. manufacturing, the percent of hours worked by individuals with less than a high school diploma falls from 33 percent to 20 percent, while the percentage of hours worked by individuals with a college degree rises from 11 percent to 18 percent.

Another means of evaluating the link to skill upgrading is to consider the occupational changes within these groups. Using CPS data for the 1970s and 1980s, Berman, Bound and Griliches (1994) show that, both in levels and changes, the white collar share of total manufacturing employment in the CPS is very similar to the nonproduction labor share generated from the ASM. The discrepancy between the white collar share and the nonproduction worker share is never greater than two percentage points over the 1970s and 1980s. It is striking that asking establishments to classify workers and asking workers to classify themselves yields such similar results.¹⁷ Given the tight link between the production/nonproduction and blue collar/white collar distinctions, Berman, Bound and Griliches (1994) show that within the white collar occupations there was an increase in the percentage of managers, professional and technical workers, and a decrease in the percent of clerical workers. They also show that within the blue collar occupations there was an increase in the

¹⁷ These results also somewhat alleviate concerns that it has become increasingly difficult for establishments to classify workers into production and nonproduction categories.

percentage of craft workers and a decrease in the percentage of operatives, laborers and service workers. Thus, the shifts among white and blue collar workers that accompanied the overall shift towards white collar workers are consistent with skill upgrading on an occupational basis.¹⁸

B. Data Sets

Our study uses three main subsets of the LRD. The first subset is the linked ASM which is an unbalanced sample of manufacturing plants for the years 1972 to 1988. This is the same data set used in Davis and Haltiwanger (1992) and Davis, Haltiwanger and Schuh (1996). The advantage of this data set is that it is a representative sample of plants in U.S. manufacturing including entering and exiting plants. The disadvantage of this data set is that the ASM is a five year rotating panel which makes longitudinal analysis across ASM panels difficult. Accordingly, the second data set we use is a subset of the linked ASM data and includes plants that appear in the ASM in all years. We refer to this data set as the balanced panel. This data set has 11,239 plants and covers approximately 38 percent of average annual manufacturing employment over the period 1972 to 1988. Figure 2 shows the nonproduction labor share in terms of employment and payroll for total manufacturing and the balanced panel over the 1972 to 1988 period. In both graphs, the basic trend in the balanced panel is virtually identical to total manufacturing. The third data set is based on the 1972, 1977, 1982 and 1987 CM's. It is a linked data set of the universe of plants which appear in each census year. This

¹⁸ Davis and Haltiwanger (1991), also using CPS data, document that much of the secular decrease in blue collar workers in the 1980s is accounted for by a sharp decrease in operatives and laborers concentrated in the 1979 to 1982 period. This latter finding is significant given that much of the overall secular increase in the nonproduction labor share depicted in Figure 1 is concentrated in the 1979 to 1982 period.

data set is similar to that used in Dunne, Roberts, and Samuelson (1989) and is used here to measure entry and exit. Appendix A provides detailed explanations and formulas for all variables used in the study.

In addition to the above data sets, we also utilize supplementary data on research and development expenditures (R&D), technology adoption, and central administrative office (CAO) employment. The R&D data come from the National Science Foundation's Annual R&D Survey which is a firm-level survey of all major R&D performers in the U.S. Using these data, we construct measures of the change in the R&D stock of the firm and match this firm-level information to the plant-level data. The information on technology use comes from the 1988 Survey of Manufacturing Technology. This survey contains data on the use of 17 manufacturing technologies including robots, local area networks, computer-automated design and flexible manufacturing cells. Both of these data sources are used to generate proxies for plant-level technology adoption that, combined with our data on workforce composition, output, and capital, are used to model capital-skill complementarity. Finally, the CAO data contain information on the employment of nonproduction labor in nonmanufacturing facilities that support the manufacturing establishments of the firm. These include headquarters, research and development laboratories, and other nonmanufacturing facilities. We use these data to examine whether changes in nonproduction labor share in manufacturing plants are related to changes in the share of firm employment in CAOs.

C. An Empirical Road Map

Our empirical analysis of the connection between plant-level technological adoption and employment dynamics proceeds in two stages. First (in sections IV-VI), we undertake a comprehensive examination of the time series evolution of the plant-level

distribution of employment changes across worker types. Second (in section VII), we investigate the connection between observable dimensions of plant-level changes in technology and the nonproduction labor share. Our motivation for the first stage is twofold. First, the aggregate (industry-level) changes in employment across worker types may provide a misleading characterization of the plant-level dynamics. Second, a change in the workforce composition at an individual plant is itself an index of a change in the state of technology at the plant. As such, we can potentially learn much about the dynamics of technology adoption by focusing on the within and between plant dynamics of workforce restructuring.

IV. Decomposition of Nonproduction Labor Share Changes

In this section, we present basic decompositions of the changes in the nonproduction labor share focusing on the relative contributions of within plant changes, between plant changes, and net entry, to the aggregate change in the nonproduction labor share. The top half of Table 1 presents the results from the decomposition given in equation (2) for the 1972-87 long difference change in nonproduction labor share using the Census of Manufactures in these two years as well as the intervening intercensal changes. The largest component is the within plant component which accounts for 43 percent of the total change in the nonproduction labor share from 1972 to 1987. The dominance of the within plant component is maintained for each of the intercensal changes, although the between plant component becomes somewhat more important in the 1977 to 1982 period (accounting for more than 25 percent). The contribution of the covariance is small and sometimes negative (less than 10 percent on average). As will become clear, underlying this relatively small covariance term is tremendous heterogeneity in the nature of the covariation

between changes in nonproduction labor share and changes in employment shares. The net entry component exhibits a somewhat different pattern in the long difference (1972-87) than in the five year changes. In the long difference, the impact of net entry is positive and substantial. In contrast, in the five year changes, the contribution of net entry is much smaller and the impact of net entry is actually negative in the 1972-1977 period. We provide a more detailed discussion of the role of entry and exit below.

The lower half of Table 1 reports summary statistics for the decomposition of annual changes in the nonproduction labor share over the period 1972-88 using the entire ASM sample and our balanced panel of continuing plants.¹⁹ The results for both the ASM sample and the balanced panel echo those from the intercensal analysis. Both the within and the between components contribute on average positively to the overall annual average change in the nonproduction labor share with the within plant component

¹⁹ A recent study by Bernard and Jensen (forthcoming) has also analyzed plant-level changes in the nonproduction labor share using a balanced panel of plants with somewhat different conclusions so it is worthwhile to discuss the sources of the differences. Their methodology and data for calculating what they denote as between and within effects are quite different from ours which makes comparisons difficult. To overcome these difficulties of comparability, we have examined the Berman, Bound and Griliches (1994) (BBG for the remainder of this footnote) type of decomposition for our balanced panel since the methodology used by Bernard and Jensen is closer in spirit to the BBG type of decomposition (although not the same given their treatment of ASM sampling weights and measurement of employment and wage shares). Using our balanced panel, the magnitudes and the time series variation in the between and within effects are very similar across the two decomposition (whether one uses an employment-based or cost-based measure). We have, however, noted some important sensitivity to the subperiods over which one examines the decomposition. Bernard and Jensen focus on two subperiods: 1973-79 and 1979-87. They argue that the between plant effect is particularly important for the latter period when one uses a cost-based measure. We find a similar result for this subperiod and for the cost-based measure as well, although unlike their results, we always find that the within plant component dominates. Over the entire 1972-88 period, we find that the annual average contribution of the between plant component to the overall change is about 1/3 using the cost-based measure and the BBG type of decomposition. For the 1979-87 period, the contribution of the between plant component is about 40 percent using the cost-share measure and the latter decomposition.

dominating. The dominance of the within plant component is especially dramatic for the analysis using the entire ASM sample. The covariance term is negative and relatively small.

Turning back to the contribution of entry and exit, at an annual frequency, the net entry effect is positive but very small. On average, entering plants tend to have a slightly lower nonproduction labor share than the initial nonproduction labor share for all plants in the year prior to entry (the average difference is -0.0038) and exiting plants have a substantially lower nonproduction labor share than the average nonproduction labor share of all plants in the year prior to exit (the average difference is -0.0119). Since entering plants have a higher nonproduction labor share than exiting plants, net entry contributes positively to the overall change. However, entering plants only constitute about 1.4 percent of employment and exiting plants only about 2.1 percent of employment at an annual frequency. These small shares imply that the overall contribution of net entry at an annual frequency is small.

In the long difference, however, entry and exit play a substantial role. To examine this issue in more depth, Table 2 reports the employment share and nonproduction labor share for entering and exiting plants for the three intercensal periods and the long difference. Three main points emerge. First, in the three intercensal periods, the employment shares of entering and exiting plants average about 12 percent. However, in the long difference, these shares rise to above 30 percent. Second, comparing the nonproduction labor share across entering and exiting plants, entering plants generally have a higher nonproduction labor share than exiting plants. In the long difference, the nonproduction labor share of entering plants exceeds that of exiting plants by $.0577$. This fact combined with the fact that entrants accumulate a substantial share over the 15

year period explains why the net entry component is relatively large in the long difference.

Third, as is true in the annual data, entering and exiting plants have a lower nonproduction labor share than continuing plants. In the case of entering plants, this may be a bit surprising since, as discussed in section II.C, it is often conjectured that entering plants will have the latest technologies and therefore should use more skilled workforces. This hypothesis neglects the fact that entrants are much smaller than continuing operations and small plants generally have a lower nonproduction labor share than large plants. In results not reported here, we compared similarly sized entering plants to continuing plants, controlling both for industry and location.²⁰ In the case of large plants (more than 1000 employees), large entrants had a higher nonproduction labor share than large incumbents. This was particularly evident in 1982 where the nonproduction labor share of large entering plants exceeded the nonproduction labor share of large continuing plants by .0456. However, smaller entrants still had a lower nonproduction labor share than smaller incumbents.

How do these findings relate to plant-level technology adoption and workforce skill? As we discussed in section II, if the aggregate nonproduction labor share change is being driven by technological upgrading of plants either through the entry-exit process or through the retooling/reorganization of existing plants, then either the within plant component and/or the net entry component should contribute heavily to the observed rise in the nonproduction labor share. Over the long run, the within and

²⁰ We estimate a simple descriptive regression of the nonproduction labor share with controls for four-digit industry, state, whether the plant is in an SMSA, and size interacted with a dummy variable indicating whether a plant is a recent entrant. The results of these regressions are available upon request.

the net entry components jointly account for 80 percent of the nonproduction labor share change. However, it is important to look more closely at the net entry component. If the difference in the nonproduction labor share between entrants and exiters arises because entry is concentrated in high nonproduction labor share industries while exit is concentrated in low nonproduction labor share industries, then changes in product demand (that induce entry and exit in this particular fashion) could generate the net-entry effect. To explore this possibility, we decomposed the change in skill due to entry-exit over the 1972 to 1987 period into a within industry component, a between industry component and a covariance term.²¹ We find that the within industry component accounts for 55 percent of the change while the between component accounts for 35 percent of the change (the remainder is the covariance component). Given this result and the generally large within plant component, it is tempting to conclude that the observed change in the aggregate nonproduction labor share reflects skill-biased technological adoption. However, as will become apparent in the succeeding sections, while this conclusion may be appropriate, the story is considerably more complicated once one considers the timing and heterogeneity in the nonproduction labor share changes across plants.

V. Cyclical Dynamics of Nonproduction Labor Share Changes

A. Between vs. Within Decomposition

²¹ Specifically, we decomposed the difference in the nonproduction labor share between entering and exiting plants in the long difference (1972 to 1987) into the within industry, between industry, and covariance components. The unit of analysis for this exercise is a four-digit industry unlike the decomposition given in (2) where the unit of analysis is the plant. Thus, for example, the within component is the difference in nonproduction labor share between entrants and exiters at the industry level weighted by the initial exiters' share.

We now turn our attention to investigating the cyclical patterns of the nonproduction labor share changes. We begin by investigating the respective contributions of the between and within plant components to the aggregate cyclical changes in the nonproduction labor share. The top panel of Figure 3 depicts the annual aggregate changes and the components of the annual changes from the decomposition for the balanced panel.²² Four results stand out. First, the dominance of the within plant component is clearly evident. Second, the nonproduction labor share change exhibits sharp increases in economic downturns in manufacturing. The sharp increases in economic downturns are only partially offset by mild decreases in the nonproduction labor share at the beginning stages of recovery (e.g., 1976 and 1984). Third, the pronounced countercyclicality of the nonproduction labor share change is primarily a within plant phenomenon. Fourth, the most pronounced positive increases in the nonproduction labor share are concentrated in the 1979 to 1982 period.

The lower panel of Figure 3 depicts the equivalent decomposition but uses the ratio of nonproduction worker wages to total wages as the measure of the nonproduction labor share. The results for the cost-share measure are very similar to the employment-share measure. The volatility of the cost-share based variable is somewhat higher than the employment-share based measure but the time series patterns are virtually identical. Again, the within plant component of the nonproduction labor share changes dominate the overall changes. The pronounced countercyclicality and the key role that the 1979-82 period play remain evident.

²² Obviously, when we consider the balanced panel we miss the role of entry and exit. However, given the small role that net entry plays in accounting for annual changes, this is not too great of a problem.

One obvious explanation for the countercyclicality of the nonproduction labor share change is the well known result that production worker employment is more cyclically sensitive than nonproduction worker employment. As discussed in section II.D, higher labor adjustment costs for nonproduction workers are likely to be relevant in this context even though we are using annual data. While this may be part of the explanation, it is far from the complete explanation of the countercyclicality. This is already evident from the asymmetry in the changes in the nonproduction labor share over the cycle. To investigate the importance of transitory changes in the nonproduction labor share more directly, we compute persistence rates for the annual within plant nonproduction labor share changes for one, two and three year horizons. The persistence rate for an individual plant over a one year horizon represents the fraction of the change from year $t-1$ to t that remains in year $t+1$.²³ The two and three year persistent measures are defined accordingly although we require the two year (three year) measure to be less than or equal to the one (two) year measure.

The average (across plants) annual one, two and three year persistence rates for the employment based nonproduction labor share changes are depicted in the top panel of Figure 4. About 60 percent of the plant-level changes persist for one year, about 45 percent persist for 2 years and about 40 percent persist for 3 years. Further, the persistence rates are relatively stable over time and do not exhibit any systematic cyclical pattern. These high and relatively stable plant-level persistence rates imply that a large fraction of the overall within plant changes reflects relatively permanent changes in the nonproduction labor share at the plant. The lower panel of Figure 4 depicts the one,

²³ The precise formulas for these persistence measures are presented in Appendix A.

two and three year persistent components of the employment-weighted, within plant, nonproduction labor share changes.²⁴ The important result that emerges is that the persistent component of the within plant changes exhibits pronounced countercyclicality. This result indicates that the countercyclicality of the within plant changes does not simply reflect the within plant shedding of production workers in recessions and then the same plant hiring production workers back in the recovery.

Even though the persistent component of the within plant changes is concentrated in downturns, the degree to which this is true varies across business cycle episodes. In the lower panel of Figure 4, the fraction of the large positive spikes in the within plant changes that persists is greater in the early 1980s than in the mid 1970s recession. This is consistent with the intercensal analysis which indicates that most of the long run changes over the 1970s and 1980s are concentrated in the 1977-82 period.

B. Job Creation and Destruction, by Worker Type

In this section, we examine job creation and job destruction by worker type. By job creation for a specific worker type, we mean the aggregate employment gains for the worker type by plants that are expanding employment of that worker type. Similarly, by job destruction, we mean the aggregate employment losses for the worker type by plants that are contracting the employment of that worker type. We convert these measures to rates by dividing the relevant flow by the average aggregate employment for that worker

²⁴ These are generated by calculating the persistence rate at the plant level for each year and then multiplying the persistence rate by the within plant change. The plant-level persistent components are then aggregated using the relevant employment shares. Thus, these measures reflect the persistent components of the first term in the decomposition in (2).

type for the prior and current period.²⁵ While the between plant job flows based on total plant-level employment have been studied extensively in the recent literature, there has been comparably little analysis of the job flows by worker type. For current purposes, the latter flows underlie not only the between plant total employment flows but also the within plant changes in the nonproduction labor share.

Table 3 presents the job creation and destruction rates by worker type at annual and five year frequencies. The rates for the annual flows are reported for the entire ASM sample (for all years) and for the balanced panel of plants.²⁶ The rates for the five year changes are based on the Census of Manufactures. Some key correlations and other summary statistics are reported for the annual tabulations in Table 4. Somewhat surprisingly, the rates of job creation and destruction for nonproduction workers are roughly on the same order of magnitude as the equivalent rates for production workers. This finding holds at both annual and five year frequencies. While the magnitudes of the rates are very similar, the magnitudes of the time series volatility are quite different across worker types. At an annual frequency, the time series standard deviations for both creation and destruction are substantially higher for production than nonproduction workers. For both worker types, job creation is procyclical, job destruction is countercyclical, the variance of job destruction is about twice that of creation, and the job reallocation (for the specific worker type) is countercyclical. Thus, the well

²⁵ Precise formulas for our measures of job creation and destruction are presented in Appendix A. Using the average of employment in the current and prior period is consistent with the methodology used by Davis, Haltiwanger and Schuh (1996). See the latter for motivation and further discussion.

²⁶ Unlike other measures reported here (e.g., those in Table 1) using the entire ASM sample, we are able to report job flows for all years. For this purpose, we exploit the methodology developed by Davis, Haltiwanger and Schuh (1996) to construct job flows for the first years of each ASM panel.

known cyclical properties of job creation and destruction for total employment hold for both production and nonproduction worker employment. In addition, the average persistence rates (constructed in an analogous manner to the persistence rates for the nonproduction labor share changes -- see Appendix A) reported in Table 4 indicate that the job creation and destruction patterns largely reflect permanent changes. The persistent job flow dynamics underlie the persistent nonproduction labor share changes depicted in Figure 4.²⁷

Within vs. Between Plant Job Reallocation

Decomposing job flows by worker type permits evaluating the contribution of within plant job reallocation to nonproduction labor share changes.²⁸ That is, the extent to which nonproduction labor share changes are undertaken via the simultaneous increase in employment of one worker type and the decrease in employment of the other worker type within the same plant. Within plant job reallocation can be measured directly at the individual plant level by summing the job creation and destruction of both worker types and then subtracting the absolute value of the net change in total employment at the plant. Along with a measure of between plant job reallocation

²⁷ Though we do not present the detailed analysis herein, we also examined the time series patterns of the persistence rates for job creation and job destruction for both production and nonproduction workers. Our basic finding is that the persistence rates are not strongly correlated with the cycle but are especially high in the 1980s. Accordingly, the high persistence rates in the 1980s underlie the especially high persistent component of the nonproduction labor share changes in the early 1980s.

²⁸ It is important to avoid confusing the between/within nonproduction labor share decomposition components with between and within plant job reallocation. A plant can exhibit large within plant changes in the nonproduction labor share with no within plant job reallocation. Similarly, a plant can contribute significantly to between plant reallocation via a large change in total employment while making no contribution to the overall change in the nonproduction labor share (e.g., if the plant has the average nonproduction labor share and makes no change in the nonproduction labor share).

(the sum of creation and destruction using total plant-level employment), total job reallocation is simply the sum of between and within plant job reallocation.²⁹

Table 5 presents estimates of between plant, within plant, and total job reallocation. Results are reported at the annual frequency for the entire ASM sample, our balanced panel and average intercensal rates from the Census of Manufactures. Annual between plant reallocation for 1972-88 is 19.2 percent of employment, within plant job reallocation is 2.7 percent, and thus total job reallocation is 21.9 percent.³⁰ Surprisingly, the relatively modest role for within plant job reallocation also holds for five year changes. Plant-level changes that involve increasing the employment of one worker type and decreasing the employment of the other worker type are relatively unimportant in accounting for total job reallocation.

One issue that Table 5 does not directly address is the importance of within plant job reallocation to changes in the nonproduction labor share. Plants can change their worker mix by either swapping one type of worker for another or by changing the scale of operations and simultaneously increasing or decreasing production and nonproduction workers at differential rates. We decomposed the change in nonproduction labor share into these two components -- the percent due to changes in scale and the percent

²⁹ Between plant job reallocation is the job reallocation measure used in Davis and Haltiwanger (1992). For the purposes of constructing measures of between, within and total reallocation, creation and destruction for all worker types and total employment are measured as a fraction of total employment (again, using the average of the current and prior period). See Appendix A for precise definition of these measures.

³⁰ While we do not report the details of the results, we have also examined the cyclical behavior of within plant job reallocation. Unlike the strongly countercyclical between plant job reallocation, the annual within plant reallocation is essentially acyclical. The acyclical nature of within plant job reallocation implies that it plays little role in the cyclical volatility of nonproduction labor share changes.

due to swapping.³¹ This decomposition shows, that on an employment weighted basis, about 30 percent of the within plant change in nonproduction labor share is accounted for by replacement of one worker type for another (swapping) and 70 percent is accounted for by scale effects. These results suggest, to the extent that the nonproduction labor share changes reflect technology adoption, this typically does not simply involve the simultaneous shedding of one type of worker and hiring of the other type of worker but rather is more likely accompanied by a change in the size of the establishment.³²

C. The Connection Between Long Run Structural Change and Cyclical Dynamics

In characterizing plant-level changes in the nonproduction labor share thus far, we have explored both long run changes and high frequency changes but we have not linked the two together. In this section, we consider some simple empirical exercises to characterize the connection between the long run structural changes and the cyclical dynamics.

We begin our analysis by examining the heterogeneity in the long run structural changes across plants. For this purpose, plants are divided into four quadrants based upon their long run changes (from 1972 to 1988) in nonproduction labor share and long

³¹ Of course, plants may be changing scale and swapping workers simultaneously (e.g., increasing one type of worker, decreasing the other type and changing scale simultaneously). Our calculations on the contribution of swapping and scale changes are thus based upon the following decomposition. For plants with employment changes of both worker types in the same direction, the swapping component is equal to zero. For plants with employment changes of each worker type in opposite directions, the swapping component is equal to the fraction of the within plant change due to pure swapping effects. The latter is calculated at the plant level from the minimum absolute value of the two employment changes.

³² The small contribution of within plant swapping of one worker type for another also implies that the within plant changes in the nonproduction worker share that we observe are *not* associated with simple relabeling of worker titles within the plant.

run changes in their share of manufacturing employment (e.g., quadrant I includes plants that increased their nonproduction labor share and their share of total manufacturing employment). For all of the exercises considered in this section we use only the balanced panel of continuing plants. Table 6 presents summary statistics of the long run changes by quadrant. The last row of the table characterizes the relative contribution of each of the quadrants to the total change based upon a modification of the decomposition given in equation (2).³³

Plants that increased both nonproduction labor share and employment share (quadrant I) accounted for almost all of the aggregate change by themselves (almost 85 percent). This result reflects the large change in nonproduction labor share and the interaction of the large change in employment share with a high initial nonproduction labor share. Plants that increased nonproduction labor share but decreased employment share (quadrant II) accounted for another 52 percent of the aggregate change. Even though quadrant II plants exhibit almost the same nonproduction labor share increase as quadrant I plants, quadrant II plants contribute less to the total because of the negative covariance (rising nonproduction labor share but falling employment share). Quadrants I and II together account for more than the total increase as their contribution is offset by plants that decreased their nonproduction labor share (quadrants III and IV). Quadrant IV's negative contribution is muted somewhat since it exhibited a large increase in employment share along with a relatively high initial nonproduction labor share. The sharply different patterns across quadrants along with the nontrivial

³³ We perform the decomposition in equation (2) in which the unit of analysis is the aggregated data by quadrant. The contribution of each quadrant to the total is given by the sum of the components of each of the three terms (between, within and covariance) for each quadrant. Given that we are using the balanced panel for this exercise, the contribution of net entry is zero.

fraction of employment in each of the quadrants highlights the tremendous heterogeneity in the nonproduction labor share changes and in the covariation between nonproduction labor share changes and changes in employment share. The finding of large scale changes (measured in terms of employment) that accompany nonproduction labor share changes helps account for the low contribution of within plant job reallocation to total reallocation reported in the previous subsection.

A primary motive for examining the behavior by quadrants is to understand the connection between the long run changes exhibited by an individual plant and the cyclical patterns of the nonproduction labor share changes for the plant. Figure 5 depicts the annual nonproduction labor share changes and employment share changes for plants in each of the quadrants. A striking feature of this figure is that quadrant I plants have positive nonproduction labor share changes in 15 of the 16 years and positive employment share increases in all years. Even with this remarkable upward consistency, quadrant I plants exhibit a pronounced countercyclicality in their nonproduction labor share changes. While there are some rough similarities in the cyclical patterns of the nonproduction labor share changes across quadrants, the experience in the early 1980s recession and subsequent recovery is very different across quadrants. Quadrants I and II concentrated much of their long run increase in the nonproduction labor share in the early 1980s recession. In contrast, quadrants III and IV concentrated much of their decrease in the nonproduction labor share in the post-1982 period. Since there are distinct groups of plants accounting for the upward vs. the downward movements and the different groups exhibit different cyclical patterns, these results reinforce the conclusion that the observed countercyclicality of the changes in the nonproduction labor share are not driven simply by plants

shedding production workers in recessions and then the same plants rehiring production workers in booms.

The different cyclical patterns across quadrants prompt us to look even deeper at the nature of the nonproduction labor share job dynamics across quadrants. Figure 6 depicts job creation and destruction rates by worker types across quadrants. Summary statistics for these job flows are reported in Table 7. The magnitudes and cyclical dynamics of the job flows vary dramatically across quadrants. In quadrant I, on average, job creation for both production and nonproduction workers substantially exceeds job destruction. The substantial increase in the nonproduction labor share and employment share in quadrant I is generated by very high rates of job creation for both types of workers with especially high rates for nonproduction workers. The volatility of destruction is about the same as that for creation for production workers in quadrant I and for nonproduction workers there is actually greater relative volatility of creation. In quadrant II, the picture is almost reversed. The substantial increase in the nonproduction labor share (and the accompanying decrease in the employment share) is generated by very high rates of job destruction for both worker types with especially high job destruction rates for production workers. The job destruction for both types of workers and especially production workers is concentrated in the early 1980s and overall there is much greater volatility of destruction as compared to creation. In quadrant III, the decrease in the nonproduction labor share and accompanying decrease in employment share is accounted for by high rates of job destruction for both worker types and enormous rates of job destruction for nonproduction workers in the 1980s.³⁴ Quadrant III also exhibits

³⁴ In evaluating these high rates, it must be emphasized that these rates
(continued...)

a pronounced asymmetry in the cyclical volatility of destruction relative to creation. In quadrant IV, the decrease in the nonproduction labor share and the accompanying increase in employment share is generated by high rates of job creation for both worker types and especially high job creation rates for production workers. The volatility of creation and destruction are about the same in quadrant IV. Overall, then, two distinct patterns emerge. First, we see that the long run changes in the nonproduction labor share are largely driven by individual plants increasing or decreasing both types of workers. Second, the overall cyclical asymmetry in the volatility of job creation and destruction (i.e., job destruction is more volatile than creation) is mostly driven by long run downsizers. For long run upsizing plants, the volatility of creation is about the same as the volatility of destruction.

D. Putting the Pieces Together

Decomposing the aggregate changes in the nonproduction labor share reveals that the cyclical patterns are dominated by within plant changes in the nonproduction labor share. Several aspects of the results point towards an important component of the latter reflecting permanent reorganizations of plant-level workforces being concentrated in downturns rather than the alternative hypothesis that the cyclical nonproduction labor share changes reflect a short run nonhomotheticity given greater cyclical flexibility of production worker employment. First, the changes are asymmetric over the cycle. Second, a large fraction of the increases in the nonproduction labor share in downturns reflect persistent changes in the within plant component of the aggregate change in the nonproduction labor share. Third, plants with a long run increase in the nonproduction labor share concentrated their increase in the recession of the early 1980s. In contrast, plants with

³⁴(...continued)

are based on the balanced panel. Recall from Table 3 that the balanced panel exhibits substantially lower rates of creation and destruction than the entire ASM panel. Thus, the rates of job destruction of nonproduction workers for quadrant III in the 1980s are astoundingly high.

a long run decrease in the nonproduction labor share concentrated their decrease after 1982.

VI. The Concentration and Timing of Plant-Level Changes in Nonproduction Labor Share

A. The Distribution of Plant-Level Changes in the Nonproduction Labor Share

Lumpy technology adoption models (like those discussed in section II) along with skill biased technological adoption imply that plants will experience large, abrupt changes in their nonproduction labor share in periods when they retool and reorganize their production process. To begin examining this hypothesis, the top panel of Figure 7 plots the distribution of annual plant-level changes in the nonproduction labor share for all plants in the ASM from 1972-88 who have non-zero employment in adjacent years.³⁵ The distribution is weighted by total employment at the plant in the initial year of each change so that each bar represents the fraction of employment in the pooled plant-year observations with a given nonproduction labor share change. Summary statistics for the distribution are reported in appendix Table A.1. There are three striking features of this figure. First, there is a very large spike at zero.³⁶ Plants with essentially no change in the nonproduction labor share in a given year constitute about 23 percent of employment. Second, while the distribution is slightly skewed to the right, reflecting the overall increase in the nonproduction labor share

³⁵ First ASM panel years, 1974, 1979 and 1984, are excluded from the tabulations used in constructing this figure because of panel rotation. Note that the distribution for the balanced panel looks essentially the same as the top panel, as does the distribution of annual nonproduction labor share changes measured by the share of nonproduction worker wages divided by total wages.

³⁶ The spike at zero is the fraction of employment with nonproduction labor share changes less than 0.005 in absolute value.

over this time period, there is tremendous heterogeneity in the distribution of nonproduction labor share changes. The relatively small mean annual change (0.0048) is generated by large annual positive gross changes (average equals 0.0231) in the nonproduction labor share by one group of plants and simultaneously large annual negative gross changes (average equals -0.0183) in the nonproduction labor share by another group of plants. Third, the distribution exhibits fat tails (excess kurtosis is very high -- see Table A.1) and the fat tails account for most of the gross changes. Plants with positive nonproduction labor share changes in excess of 0.05 (evaluated at the average nonproduction labor share this constitutes almost a 20 percent change) in a given year account for more than 70 percent of the overall aggregate positive nonproduction labor share change, while plants with a negative nonproduction labor share change of less than -0.05 account for more than 70 percent of the overall negative nonproduction labor share change.

The lower panel of Figure 7 depicts the distribution of long differences in the plant-level nonproduction labor share for all plants with positive employment in both the 1972 and 1987 Census of Manufactures. Not surprisingly, a much smaller percent of employment is found in plants with little or no change in the nonproduction labor share (about 6 percent). The most striking feature again is the important role of both positive and negative changes in accounting for the overall change.

One possible explanation for the changes in the nonproduction labor shares, and in particular the large negative tails observed in both panels in Figure 7, is that firms are shifting their nonproduction employment between their manufacturing facilities and into central administrative offices (CAOs). We explore this possibility by examining changes in nonproduction employment in both the manufacturing facilities

and CAOs of the same firms.³⁷ Overall, there is a negative, but quite weak, correlation between changes in nonproduction labor share in manufacturing facilities and the change in employment share of CAOs (- .06). In addition, we examine whether firms with large negative changes (less than -.05) in nonproduction labor share in their manufacturing facilities had large offsetting increases in CAO employment. If one factors in the change in CAO employment into the overall change in nonproduction labor share for firms experiencing large declines in nonproduction labor share, the impact on the overall change is quite modest (the employment-weighted average decline for such firms is -0.12 without the CAO change and -0.09 with the CAO change). Our conclusion is that the large negative change in nonproduction labor share cannot be accounted for by firms simply shifting nonproduction labor from manufacturing facilities to CAOs.

B. The Timing of Large Plant-Level Nonproduction Labor Share Changes

While the findings depicted in Figure 7 are consistent with lumpy adoption models, there are a number of alternative explanations for this lumpiness. As discussed in section II, fixed costs of adoption and lumpy technological improvements are two (not mutually exclusive) possible explanations. Yet another explanation is an extreme form of short run nonhomotheticity associated with plants with large changes in the scale of operations (induced by factors other than changes in technology). One way to help distinguish among these hypotheses is to examine the timing of large scale changes in the nonproduction labor

³⁷ The analysis we undertake examines all multi-plant firms. Note, that the vast majority of manufacturing establishments are single plant operations with no separate headquarters facilities. In addition, most multi-plant operations also do not have separate headquarter facilities. Roughly, 90 percent of all manufacturing establishments are owned by firms with no CAOs. However, those firms with CAOs are quite large and account for approximately 55 percent of total manufacturing employment over the 1977 to 1987 period.

share at the plant. As discussed in section II, one prediction of lumpy technology adjustment models based upon a vintage capital specification is that the probability of retooling should be increasing in the time since the last retooling. We investigate this prediction in an indirect manner by examining whether the probability of having a large nonproduction labor share change is increasing in the time since the last large nonproduction labor share change.

For this purpose, we divide plants into two groups in a manner similar to our quadrant analysis in section V. Plants with long run increases in the nonproduction labor share are in one group and plants with long run decreases in the nonproduction labor share are in the second group. The motivation for this grouping is that plants with positive long run changes in the nonproduction labor share likely adopted different types of technology than plants with negative long run changes. For plants with positive long run changes, we define a nonproduction labor share spike to be equal to one (zero otherwise) in any year in which the change in the nonproduction labor share exceeds 0.05. For plants with negative long run changes, we define a nonproduction labor share spike to be equal to one (zero otherwise) in any year in which the change in the nonproduction labor share is less than -0.05.³⁸

We estimate variants of the following simple specification:

$$h_{it} = \alpha_i + \sum_k \alpha_k D_{kit} + \alpha_c * CYC_{it} + \epsilon_{it} \quad (3)$$

where $h_{it} = 1$ ($=0$) if plant i has (does not have) a nonproduction labor share spike in period t , α_i represents a possible plant-fixed effect,

³⁸ These threshold values for spikes are arbitrary but from Figure 7 clearly represent very large changes relative to the typical change. We have considered alternative thresholds (0.03 and 0.10) and have obtained very similar results for the hazards.

D_{kit} represent "age" dummies reflecting the number of years since the prior nonproduction labor share spike (e.g., $D_{kit} = 1$ if the plant last had a nonproduction labor share spike k years ago, zero otherwise), and CYC_{it} represents a cyclical indicator. The indicator we use for this purpose is a downstream demand indicator specific to the four-digit industry to which the plant is assigned (this measure is developed in Bartelsman, Caballero, and Lyons, 1994). The downstream indicator is the change in an index of activity of other industries and service sectors which purchase output from the industry in question.³⁹

We consider two alternative specifications of (3). First, we estimate (3) via OLS without any fixed effects. We denote the results from this estimation as the Kaplan-Meier estimates since, in the absence of controlling for the cyclical indicator, the coefficients on the age dummies reflect the simple empirical hazard. Second, we attempt to control for unobserved heterogeneity that lead some plants to have systematically high probabilities of spikes while others have low probabilities of spikes. We do this because failure to control for such systematic unobserved heterogeneity will bias the hazard downwards. One method of addressing this problem is to estimate a first difference specification of (3) in order to eliminate fixed effects. However, by taking first differences, we induce a correlation between the difference of the age dummies and the difference in the errors. To overcome this problem, we estimate the first difference specification with instrumental variables where we use twice (and greater) lagged dependent

³⁹ The use of this downstream demand cyclical indicator is motivated by the arguments made by Shea (1993) that downstream demand is arguably a good instrument for industry-level demand variation. Of course, as Shea argues, the downstream demand indicator is a better exogenous instrument when the material share of the output from the upstream industry in the total costs of the downstream industry is relatively low.

variables and other twice (and greater) lagged plant-specific variables as instruments.⁴⁰

For this purpose, we exploit the balanced panel data from 1972-88. In order to construct the age dummies for several periods, we commence the estimation of (3) in 1978 and consider the variation across 5 age groups. Specifically, the groups are plants that last had a large nonproduction labor share change in the prior year (age=1), two years prior (age=2), three years prior (age=3), four years prior (age=4) and five or more years prior (age=5 -- the omitted group in the estimation). For plants with positive long run nonproduction labor share changes, the average number of spikes is 1.75. The equivalent mean for plants with long run negative nonproduction labor share changes is 1.70.

The results from this estimation are reported in the panels of Table 8 and the hazards for plants with long run positive nonproduction labor share changes are depicted in Figure 8. In the latter, the line labeled "boom" reflects the hazard that results by allowing the industry-specific cyclical indicator to be one standard deviation above its mean for each plant and then averaging across plants. The line labeled "recession" is the equivalent one standard negative deviation case. For the Kaplan-Meier estimates, the estimation yields not only the slope and cyclical sensitivity but also the appropriate level (the intercept) for the hazard. For the first difference specification, we obtain only the slope and cyclical sensitivity so some caution needs to be used in comparing the results.

The results for both the positive and negative long run changes are qualitatively similar so we focus our discussion on the positive

⁴⁰ In addition to the twice (and greater) lagged dependent variables the instrument list includes twice lagged capital growth (equipment and structures), output growth, indicators of changes in ownership and industry, and the downstream industry demand indicator. The capital and output growth are interacted with two digit industry dummies.

long run changes. For the Kaplan-Meier estimates, we observe a hazard that rises initially but ultimately decreases. Although the slope is not steep, the first difference hazards are upward sloping throughout. Thus, even this crude control for selection effects yield results that are qualitatively consistent with the prediction of an increasing hazard in a lumpy adjustment model. Plants are more likely to experience a large positive spike in their nonproduction labor share the longer it has been since the prior positive spike.

The results also indicate that positive spikes in the nonproduction labor share are countercyclical. This finding echoes the results reported in the previous sections regarding the pronounced countercyclicalities of within plant nonproduction labor share changes. The additional result here is that this countercyclicality is driven in part by the countercyclicalities of spikes in the nonproduction labor share.

The large spike at zero in the distribution of nonproduction labor share changes, the fat tails of the distribution, and the rising hazard, are all consistent with lumpy adjustment models. Further, the countercyclicalities of the positive nonproduction labor share spikes is consistent with the idea that recessions may be good times to reorganize and retool since the opportunity cost of foregone output is low at such times.⁴¹ Nevertheless, these results on the distribution of nonproduction labor share changes raise a variety of questions. The large positive and negative gross changes in the nonproduction labor share that are dominated by large (absolute) spikes in the nonproduction labor share raises questions about the nature of bias in technical

⁴¹ Although the results for plants with negative long run nonproduction labor share changes indicate that the latter plants do not concentrate their changes in recessions. Based upon the magnitudes of the coefficients on the cyclical indicator, the procyclicality of the latter group is less strong than the countercyclicalities of the plants with long run positive changes in the nonproduction labor share.

change. One obvious specific question is: What is occurring at plants experiencing large decreases in their nonproduction labor share? One possible explanation is that the technology-skill complementarity varies (in sign and magnitude) across different types of technologies. If technology that is primarily used in production is complementary with unskilled workers then plants which adopt this type of technology will decrease their nonproduction labor share.⁴² The challenge, then, is to find observable plant-level variables that can account for both the left and right tails of the distribution of nonproduction labor share changes. We turn to this challenge in the next section.

VII. The Relationship Between Observable Measures of Technology Adoption and the Nonproduction Labor Share

We now turn to investigating the connection between observable measures of technology adoption and changes in the nonproduction labor share at the plant level. Our empirical methodology is similar in spirit to Bartel and Lichtenberg (1987) and Berman, Bound, and Griliches (1994). We start by assuming that firms minimize a cost function that contains both variable and quasi-fixed inputs. In our case, we specify a translog cost function with two variable factors of production (skilled and unskilled workers) and a number of possible quasi-fixed factors. Applying Shepherd's Lemma and taking first differences yields:

$$\Delta M_{it} = \alpha_i + \alpha_0 \Delta \ln(w_t^s/w_t^u) + \sum_j \alpha_j \Delta Z_{ijt} + \alpha_y \Delta \ln(Y_{it}) + \epsilon_{it} \quad (4)$$

⁴² We explore this idea further in section VII when we look directly at the effect of different types of plant-level technology adoption on the nonproduction labor share in the plant. See Dunne and Troske (1995) for further evidence on this issue.

where in this case M_{it} is measured as the share of nonproduction workers wages in the total wage bill for plant i in period t .⁴³ The Z_{ijt} 's include the quasi-fixed factors. Real output, Y_{it} , is included to capture possible nonhomotheticity. μ_t captures a common time effect such as the common unobserved component of the bias in technological change and η_{it} captures unobserved idiosyncratic biased technical change (including unobservable Z_{ijt} 's). For a particular quasi-fixed factor Z , the coefficient on the Z represents the Z -skill complementarity. In what follows, we estimate (4) using both the change in the cost-share based measure of the nonproduction labor share and the change in the employment-share based measure of the nonproduction labor share as the dependent variable. While the latter does not emerge from the translog cost functional, it is consistent with the general form of the nonproduction labor share characterized in equation (1).

There are a number of issues that must be addressed before estimating the model. First, we must determine the variables in Z_{ijt} . In all of our specifications, Z_{ijt} includes plant-level log of capital equipment (beginning of period) and log of capital structures (beginning of period). In an attempt to capture any change in organizational structure, we also control for whether the plant changes ownership or industry in the period in question. In some specifications, we include a firm-level measure of the log of the stock of R&D as well as direct plant-level measures of the type of technology adopted at the plant from the 1988 Survey of Manufacturing Technology. In both of the latter cases, inclusion of such variables restricts the sample and/or the time

⁴³ This specification is formally derived in Brown and Christensen (1981) and Bartel and Lichtenberg (1987) amongst others. Note also, that equation (4) closely resembles the relative factor demand equation (equation (1)) derived in section II.

period of the analysis in important ways (details below) and thus these variables are not included in all specifications.⁴⁴

A second issue is the treatment of relative wages. There is likely little directly measurable exogenous cross sectional variation in relative wages that we can exploit for this purpose. The problem is that wage changes for a given worker type are confounded with unobserved labor quality differences. Extensive analysis of the plant-level wage variation (e.g., Davis and Haltiwanger, 1991; Troske, 1994; Doms, Dunne and Troske, 1995) indicates that much of the between and within plant wage variation by worker type reflects differences in labor quality between and within plants. Even though we cannot directly measure changes in relative wages, to the extent that relative wages for skilled workers are equalized across the economy, such economy wide variation in relative wages will be captured through our time effects. We attempt to improve upon this by interacting the time effects with region dummies. The region-time effects are included in all of the specifications discussed below.

A third issue is endogeneity bias for the various measures of the Z_{ijt} 's. Abstracting from measurement error for the moment, the error term in (4) reflects the current period idiosyncratic shock to biased technological change.⁴⁵ This shock is likely to be correlated with the current growth rate of output so there are clear endogeneity problems here. Since the change in the log of capital in (4) represents the change from the beginning of the prior period to the beginning of the current period, at annual frequencies it is reasonable to argue that last period's investment is not correlated with the current period

⁴⁴ The details of the construction of these variables is discussed in Appendix A.

⁴⁵ Since our specification is in first differences, this is really the innovation to the current shock.

technology shock. However, in what follows, we consider not only one year differences but also three year and long differences (1978 to 1988). In the latter cases, it is much more likely that the cumulative changes in the capital stocks are correlated with the cumulative technology shocks.

A related issue is measurement error in the Z_{ijt} 's. Following Griliches and Hausman (1986), part of the motivation for considering three year and long differences is to mitigate the effects of measurement error (in addition, the longer differences are interesting in their own right). However, as just discussed, this in turn generates problems with endogeneity. Putting these concerns together, in what follows we estimate our specifications both via OLS (as a benchmark and to compare to the results of the recent studies that have used OLS in related settings -- e.g., Berman, Bound, and Griliches, 1994; Bernard and Jensen, forthcoming; Goldin and Katz, 1996) and via an instrumental variables procedure where we treat the changes in equipment, structures and output as endogenous (as well as other relevant Z measures as appropriate). We instrument for these variables using appropriately lagged plant-level measures. Given our concerns about measurement error and again following Griliches and Hausman (1986), for the difference between $t-k$ and t we use plant level variables from period $t-k-1$ (and earlier) as instruments.⁴⁶ In all of our specifications the initial $t-k$

⁴⁶ The precise instrument list for the difference between $t-k$ and t are plant level variables dated $t-k-1$ (and earlier) including linear, quadratic, cubic and quartics in the log of output, the log of capital equipment, and the log of capital structures as well as quadratic, cubic and quartic terms in the growth rates of each of these lagged variables. The motivation for including the nonlinear terms is based on the recent literature indicating that plant-level dynamics in employment and capital exhibits significant nonlinearities (see, e.g., Caballero, Engel and Haltiwanger, 1995a, 1995b). We also include the contemporaneous and lagged downstream cyclical indicator used in section VI as additional instruments in all specifications. For alternative specifications including additional Z_{ijt} 's (e.g., R&D), the appropriate $t-k-1$ lags (and earlier) for these additional variables are included as instruments as well.

is 1978 (e.g., the first one year difference is 1978-1979, the long difference is 1978 to 1988). Since much of the within plant changes occur over the 1978 to 1988 period, we are able to include the most interesting period of variation in our sample period. We estimate the instrumental variables specification via GMM to improve efficiency and to generate consistent standard errors.

In all cases we prefer to estimate the model separately for plants in each two-digit industry. An obvious advantage of using plant-level data is the ability to let parameters vary across observable dimensions of differences in technology such as industry. Even in the specifications estimated separately for two-digit industry, we interact three-digit industry effects with the time effects to capture differences across industries within two-digit industries.

A. Base Specification -- The Contribution of Capital-Skill Complementarity

Table 9 presents estimated coefficients for the pooled OLS and IV-GMM specifications. While we ultimately draw our inferences from results that are estimated separately for plants in each two-digit industry, we present the results in Table 9 to provide a benchmark for comparison across alternative specifications and for comparability with the existing literature. Several aspects of these results deserve mention. First, there is clear evidence of capital-skill complementarity with the complementarity between skill and equipment generally stronger than the complementarity between skill and structures. Second, the IV-GMM results generally yield higher absolute magnitudes of coefficients especially in the one year difference specifications. This latter finding is consistent with measurement error in our Z_{ijt} 's. Third, the OLS results yield generally higher magnitudes of the capital-skill complementarity in the long differences

as opposed to the one year differences which is again consistent with measurement error in our right hand side variables. This pattern is less pervasive in the IV-GMM results which makes sense if one interprets the one year differences with IV-GMM as being purged of measurement error.⁴⁷ Fourth, we find clear evidence of short run negative non-homotheticity in the one and three year difference results. However, this negative nonhomotheticity diminishes in the longer differences. Fifth, the change in ownership and change in industry affiliation variables exhibit somewhat erratic patterns in these specifications. However, in the long run it appears that both events are associated with increases in the nonproduction labor share.⁴⁸ Sixth, the low overall R-squared's indicate the dominance of unobserved factors in accounting for the pooled cross-sectional and time series variation. Given the inclusion of both year effects and year effects interacted with industry and region, it is clear that there is tremendous heterogeneity in the nonproduction labor share changes in a given year across plants in the same industry or region. Finally, the results for the employment-share based specification and the cost-share based specification are qualitatively very similar. This pattern holds throughout the remainder of the subsequent analysis and for the sake of brevity we focus on the employment-share based results for the remainder of the paper.

⁴⁷ Our IV approach is much less successful in our long differences specification than in our one or three year specifications. This is because, given the limited time dimension of our data, we are unable to construct instruments which accurately predict the long run changes in the endogenous variables.

⁴⁸ The industry change and ownership change indicators are much better predictors of changes in the absolute nonproduction labor share than the nonproduction labor share itself. We estimated alternative specifications of (4) with the same right hand side variables but with the dependent variable being the absolute change in the nonproduction labor share. The coefficients on these two indicators (results not reported) increased substantially in magnitude and statistical significance.

As we stated above, one advantage of the microeconomic data is that it allows us to estimate the above model separately for each two-digit industry. Table 10 reports these results from the one year difference specification using IV-GMM. We report coefficient estimates as well as the employment-weighted average plant-level growth rates for net investment and output. It is evident that there are tremendous differences across two-digit industries in the nature and magnitude of the equipment-skill and the structures-skill complementarity and in the short run nonhomotheticity. For example, the equipment coefficient varies from -0.013 in Printing to 0.068 in Instruments. While the industries with negative equipment-skill complementarity are imprecisely estimated, there is nevertheless a wide range of values across industries. There are also striking differences in the growth rates of net investment and output across industries. Given the large differences in the coefficients and the large differences in the growth rates of these variables across industries, there are clearly problems in using the pooled estimates (which constrain the coefficients to be the same across industries) to make inferences about the contribution of the relevant variables to the variation in nonproduction labor share changes.⁴⁹

Our main objective in this exercise is to understand the contribution of our observable factors to the time series and cross sectional variation in the within plant changes in the nonproduction labor share. For this purpose, we present a simplified version of the full distribution accounting framework used by Juhn, Murphy and Pierce (1993). That is, we decompose the employment-weighted distribution of nonproduction labor share changes into the changes in the distribution accounted for by the observables and the changes in the distribution

⁴⁹ Thus, studies based upon aggregate (industry-level) data that constrain the coefficients to be the same across industries are potentially misleading.

accounted for by the unobservables.⁵⁰ By observables, we mean the predicted plant-level changes in the nonproduction labor share in each year based upon all of the explanatory variables except the common year effect. By unobservables, we mean the plant-level changes in the nonproduction labor share in each year accounted for by the estimated common year effect plus the estimated plant-level residual. We treat the estimated common year effect as unobservable since it represents the unexplained common bias in technological change.

Decomposing annual employment-weighted mean changes in the nonproduction labor share simply involves computing the annual employment-weighted mean change in the nonproduction labor share from the observables and the annual employment-weighted mean change in the nonproduction labor share from the unobservables. However, we expand the standard mean decomposition (or the standard variance decomposition) by decomposing the entire distribution of nonproduction labor share changes into the contribution of the observable and unobservable factors. Consider, for example, the decomposition of the difference between the 95th and 5th percentile of the employment-weighted nonproduction labor share change distribution. In a given year, we measure the contribution of the observables as the difference between the 95th and 5th percentile of the distribution of plant-level predicted changes in the nonproduction labor share. The marginal contribution of the unobservables is then given by the difference between the actual 95-5 and predicted 95-5 percentile differences.⁵¹

⁵⁰ The relevant employment weights in each year are the initial (prior year) employment shares. Thus, this exercise decomposes the within plant component of equation (2).

⁵¹ Unlike the decomposition of the mean change or the analogous variance decomposition, the ordering of the decomposition potentially matters in this context. That is, one could start with the contribution of the unobservables to the 95-5 difference and then consider the marginal contribution of the observables.

Figure 9 depicts the decomposition of the annual mean changes in the nonproduction labor share as well as the decompositions of the annual 95th-5th percentile difference, the annual 95th-50th percentile difference and the annual 50th-5th percentile difference. In each case, we present the statistic from the actual distribution, the contribution to this statistic by the observables and the marginal contribution to this statistic by the unobservables. For the contribution of the observables, we also present the contribution from the observables minus the effect of plant-level changes in output. For this exercise we use the estimates from the IV-GMM one year difference specification estimated separately for plants in each two-digit industry.

The upper left panel of Figure 9 depicts the decomposition of the employment-weighted annual mean nonproduction labor share changes. Over the period, observables account for about 40 percent of the average annual change in the nonproduction labor share. Excluding the contribution of plant-level output changes, the capital component of the observables accounts for about 30 percent of the average annual change.

While the observables account for a sizeable fraction of the average annual change, the observables (and in particular capital-skill complementarity effects) do not track the timing of the nonproduction labor share changes well. For example, in 1981 the observables account for only 27 percent of the sharp increase in the nonproduction labor share and capital-skill complementarity accounts for only 10 percent of the increase. More generally, capital-skill complementarity predicts a positive nonproduction labor share change in all years and relatively modest time variation.⁵²

⁵² In examining the precise timing of the contribution of capital-skill complementarity in Figure 9 it is important to recall that it is beginning of period capital stock that affects the current period nonproduction labor

(continued...)

Changes in output account for some of the cyclical variation in the nonproduction labor share consistent with short run nonhomotheticity induced (at least in part) by the greater cyclical flexibility of production workers. Nevertheless, unobservable factors dominate the cyclical variation. Further, unobservable factors account for most of the secular increase in the nonproduction labor share. Putting the two results together implies that unobservable factors are generating most of the long run increase in the nonproduction labor share and that the long run changes generated by these unobservable factors are concentrated in recessions. It should not be surprising that capital-skill complementarity cannot account for these cyclical patterns. There is an inherent tension in trying to account for sharp increases in nonproduction labor share changes in downturns with capital-skill complementarity since net investment is procyclical.

Turning now to the decompositions of the percentile differences, several interesting patterns emerge. First, the actual 95-5 percentile difference indicates a secular increase in the dispersion of nonproduction labor share changes over the 1980s. This secular increase is again accounted for entirely by unobservable factors. In addition to the secular increase in dispersion, there is an increase in dispersion in the cyclical downturn in the early 1980s that is again primarily accounted for by the unobservables. Looking at the 50-5 and the 95-50 differentials in the nonproduction labor share change distribution shows that there is a greater secular increase in the dispersion in the left tail of the distribution (the 50-5 difference) but greater cyclical volatility in the dispersion in the right tail of the distribution (the

⁵²(...continued)

share. Accordingly, in first differences it is the change in the capital stock in the prior period that affects the change in the nonproduction labor share from the prior to the current period. Thus, for example the equipment investment boom in 1985 yields an increase in the nonproduction labor share in 1986.

95-50 difference). Observables account for very little of either the secular or cyclical variation in the left tail or the secular change in the right tail. However, observables account for a relatively larger fraction of the cyclical volatility in the right tail.

B. The Contribution of R&D Investment

We now consider the additional contribution of net changes in the stock of R&D capital at the plant level. Essentially, we follow a methodology similar to that used by Adams and Jaffe (1994) by building stocks of R&D at the firm-level using a perpetual inventory method (details discussed in Appendix A). The firm level R&D stocks are matched to individual plants using the firm identifiers in our plant-level data set. We restrict our sample to those plants for which we have a continuous series on R&D, which reduces our sample by about 40 percent.

Table 11 reports the coefficient estimates with the inclusion of the net investment in R&D for the pooled OLS, IV-GMM and the employment-weighted, two-digit industry, for the both the one year and three year difference specifications.⁵³ The results for net investment of equipment and structures are largely similar to those reported in Tables 9 and 10. We find evidence of significant R&D-skill complementarity in the IV-GMM results for both the one and three year difference specifications.

The decomposition of the employment-weighted nonproduction labor share distribution into observable and unobservable components when we

⁵³ We treat changes in the R&D stock as endogenous in our IV-GMM specifications. Once again, we do not report the long difference estimations because, given our small sample size, we do not feel we can use IV to estimate the model.

include R&D is presented in Figure 10.⁵⁴ The results are qualitatively similar to those presented in Figure 9. Observables account for 45 percent of the average annual change and observables absent the output growth effect account for 42 percent of the average annual change. Including net R&D investment does not help in accounting for aggregate cyclical changes in the nonproduction labor share or to the secular increase in the dispersion in the nonproduction labor share changes. This finding is again not surprising given the essentially acyclical nature of R&D investment. There is a modest improvement in the contribution of the observables to the cyclical volatility in dispersion (especially for the right tail). In short, however, the main conclusion of the prior section still holds: unobservable factors account for most of the nonproduction labor share increases in cyclical downturns and the cyclical increases generated by unobservables account for most of the long run secular increase in the nonproduction labor share.

C. The Contribution of Indicators of Adoption of Advanced Technology

We now consider the influence of direct measures of the type of technology adopted at the plant. For this purpose, we exploit the 1988 Survey of Manufacturing Technology (SMT). As discussed in section III, the 1988 SMT surveyed 10,590 plants in SIC industries 34 to 38 on their use of seventeen recent manufacturing innovations. The innovations include such technologies as robotics, computers on the factory floor, local area networks, computer-aided design and automatic sensors. We match the data from the SMT with our balanced panel of continuing plants which results in a sample of 1820 plants. Following Dunne and Troske (1995), we divide the technologies into two main groups -- information technologies including computer automated design (CAD), computers and

⁵⁴ The results in Figure 10 are based on the one-year difference IV-GMM specification estimated separately for plants in each two-digit industry.

lans, and production technologies including flexible manufacturing cells, lasers, and pick and place robots. The distinction between information technologies and production technologies is that the former are technologies that aid in the management and design of products and the production process while the latter aid directly in production.

Our analysis exploiting the SMT data has a number of important limitations. First, we only have the information at one point in time for our balanced panel (1988). Thus, we cannot exploit the precise timing of adoption. To get around this problem, we use the SMT variables in the long difference specification and presume that if a particular technology is present at an individual plant then the adoption of the technology occurred between 1978 to 1988. Second, we estimate the specification by OLS and IV-GMM both with and without instrumenting for technology adoption using the same instrument set as in the base specification. Since the need (at least for measurement error reasons) to instrument in the long difference specification is limited and because we do not have any reasonable instruments for the SMT variables, we have more confidence in the OLS results in this case and focus our discussion on these results. Third, the results are restricted to a narrow set of industries (SIC 34 to 38) which restricts the sample size and the relevance of the analysis for total manufacturing and makes comparisons to the preceding analyses more tenuous.

The results for the long difference specification including the SMT variables are reported in Table 12. For both the information and production technologies, we break plants into two groups based upon whether the plant is using at least one of the relevant technologies. The results for capital-skill-complementarity and nonhomotheticity are quite similar to those generated in the base case. We find that adoption of information technologies yields a positive impact on the

change in the nonproduction labor share while the adoption of production technologies generates a negative impact on the change in the nonproduction labor share.⁵⁵ Since these results reflect the influence of the type of technology adopted after already controlling for output and capital growth, these results do not necessarily imply that plants that have adopted information technologies exhibit increases in their nonproduction labor share or that plants that have adopted production technologies exhibit decreases in the nonproduction labor share. Important in this regard is that plants that have adopted either type of advanced technology generally have much higher net equipment investment rates. For example, the employment-weighted mean net investment rate for plants that have adopted advanced production technologies is 0.48 while plants that have not adopted such technologies have a rate of 0.297 (both rates over the 10 year period).⁵⁶

⁵⁵The results reported in Table 12 should be viewed with some caution. First, the sample of plants is dominated by very large producers and 91 percent of plants in the sample use information technologies. Second, Doms, Dunne and Troske (1995), using a larger sample of plants, a more recent SMT, and employing the number of technologies used as their main index of technology adoption, find little correlation between technology use in 1993 and the change in nonproduction labor share over the 1977 to 1992 period. In results not reported here, we estimated the nonproduction labor share regression using a count-based measure and found little correlation between the number of technologies used and the change in nonproduction labor share. Finally, Doms, Dunne, and Troske also estimate a model similar to that reported in Table 12. They find a negative relationship between production technology use and the change in nonproduction labor share (as we do), but find a much weaker positive correlation between information technology use and changes in nonproduction labor share than we find in our more restricted sample.

⁵⁶ The employment-weighted aggregate change in the nonproduction labor share for plants that adopted information technologies is 0.0353 while the aggregate change for those that have not is 0.0139. The predicted aggregate change for those that have adopted information technologies is 0.0319 while the predicted aggregate change for the other plants is 0.0256. The employment-weighted aggregate change in the nonproduction labor share for plants that adopted advanced production technologies is 0.0333 while the aggregate change for those that have not is 0.0434. The predicted aggregate change for those that have adopted production technologies is 0.0293 while the predicted aggregate change for the other plants is 0.0461. Thus, we observe that on average plants in all four of these groups exhibit substantial

(continued...)

In terms of the contribution of the observables to the long run nonproduction labor share change, in a manner consistent with our above annual analysis we treat the constant as an unobservable. The actual employment-weighted average nonproduction labor share change for the 1820 plants over the 1978-88 period is .035. The predicted change in nonproduction labor share when we include our SMT variables is .017, so the observables account for about one-half of the change in the nonproduction labor share in this sample of plants. However, even after controlling for whether or not a plant adopts new technologies, unobservables still account for 50 percent of the secular employment-weighted average change in the nonproduction labor share.

VIII. Concluding Remarks

Our main empirical findings are summarized as follows:

1. Aggregate changes in the nonproduction labor share at both annual and longer frequencies are dominated by within plant changes in the nonproduction labor share. The contribution of between plant changes in the nonproduction labor share is relatively modest peaking (around 25 percent) in the 1977-82 period. Net entry reinforces the effects of the within plant changes since the nonproduction labor share differential between entering and exiting plants is similar to the within plant changes in the nonproduction labor share that we observe for continuing plants. Taken together, the contribution of within plant changes and net entry jointly account for more than 80 percent of the long run change in the nonproduction labor share.

(...continued)

increases in their nonproduction labor share but particularly high changes for plants that have adopted information technologies and have not adopted production technologies.

2. The distribution of annual within plant changes exhibits a spike at zero, tremendous heterogeneity and fat left and right tails. The aggregate net change in the nonproduction labor share is driven by large gross positive and negative changes in the nonproduction labor share that are an order of magnitude larger than the net change. Most of the large gross changes are accounted for by plants with very large (in excess of .05 in absolute value annually) lumpy changes. The probability of a plant experiencing a large change in the nonproduction labor share is increasing in the time since the previous large change in the nonproduction labor share.

3. Within plant job reallocation is small in magnitude and accounts for only a small fraction of total job reallocation (between plus within plant job reallocation). Additionally, most changes in the within plant nonproduction labor share are driven by individual plants either simultaneously increasing both skilled and unskilled workers or simultaneously decreasing both types of workers. Annually, the rate of within plant job reallocation is only 2.7 percent. This compares to a between plant job reallocation rate of about 20 percent. Most changes in within plant nonproduction worker shares are accounted for by plants changing scale but disproportionately via employment of one type of worker. Plants that exhibit long run secular increases in the nonproduction labor share accomplished through long run upsizing account for most of the long run secular increase in the nonproduction labor share. This group of plants accomplish this with large rates of job creation for both skilled and unskilled workers and especially high rates of job creation for skilled workers. Another group of plants that increase their nonproduction labor share through long run downsizing also contribute significantly to the long run secular increase in the nonproduction labor share. This latter group of plants accomplish this

with high rates of job destruction for both worker types and especially high job destruction rates for unskilled workers.

4. The long run secular increases in the nonproduction labor share are concentrated in recessions, particularly the recession in the early 1980s. The aggregate nonproduction labor share rises sharply in economic downturns and falls only mildly in economic recoveries. The persistent component of within plant nonproduction labor share changes increases sharply in recessions. Plants with long run increases in their nonproduction labor share concentrated much of their increase in the recession in the early 1980s. Plants with long run decreases in the nonproduction labor share account concentrated much of their decrease after 1982.

5. Observable indicators of changes in technology account for a significant fraction of the long run secular increase in the average nonproduction labor share but account for little of the cyclical variation. Unobservable factors account for most of the long run secular increase and the changes generated by these unobservable factors account for most of the cyclical variation in the aggregate nonproduction labor share. Putting these results together, the factors that are generating the long run secular increase in the nonproduction labor share are closely linked to the factors generating the cyclical increases in the nonproduction labor share in economic downturns.

Capital-skill complementarity (including detailed measures of the type of advanced technologies adopted) and R&D-skill complementarity contribute significantly to the employment-weighted average secular change in the nonproduction labor share. However, capital investment and R&D investment account for little of the cyclical variation in the average within plant nonproduction labor share changes. Even after controlling for plant-level net investment of physical capital, net investment in R&D capital and output growth, the residual average within

plant changes in the nonproduction labor share exhibit significant countercyclicality. Further, the increases in the average within plant nonproduction labor share generated by unobservable factors in downturns are much larger than the modest decreases in recoveries, so that the cumulative effect is positive. Indeed, more than half of the long run secular increase in the average nonproduction labor share is accounted for by these unobservables.

These results, in general, point to the conclusion that the aggregate change in the nonproduction labor share reflects some form of technical change (broadly defined) inducing individual continuing plants to retool and reorganize and new more skill intensive plants to displace exiting less skill intensive plants. The results are consistent with models that predict that retooling and reorganization will be lumpy and that the latter will imply similar lumpiness in nonproduction labor share changes. Further, the results are consistent with models that predict that retooling and reorganization of the workforce will be concentrated in recessions since the opportunity cost of the disruption from undertaking retooling/reorganization is low during recessions.

Viewed from the broad perspective of changes in the manufacturing workforce, our results indicate that there have been striking changes within individual manufacturing establishments in the type of workers used at operating production establishments. While on average the shift has been towards nonproduction workers, the underlying plant level changes reflect tremendous heterogeneity in the within plant changes in the structure of employment. Our results do not provide much support for simple product demand and related trade explanations for the change in the mix of jobs in U.S. manufacturing. Instead, our results are consistent with the view that individual plants have fundamentally changed the way they produce goods in terms of the mix of workers employed at production establishments.

There are, however, more questions raised than answers provided by our analysis. The dominant role that unobservable factors play in accounting for most of the secular and cyclical variation in the nonproduction labor share raises questions about what these unobservable factors represent. In a similar vein, the tremendous heterogeneity in within plant changes in the nonproduction worker share (with large positive and negative gross within plant changes) raise questions about the nature of the bias in technical change. We have argued that one label to put on the factors generating these patterns of plant-level changes is organizational capital. Viewed from this perspective, our results indicate that organizational capital has on average been skill biased over this period and that permanent reorganizations of the production process at the plant are concentrated in recessions. Organizational capital need not be skill biased in general and indeed we find that unobservable factors account for most of the tremendous cross sectional variation in plant-level nonproduction labor share changes.

The dominant role of unobservables is not inherently in conflict with the recent interpretations from related aggregate analyses (e.g., Berman, Bound and Griliches, 1994) that the aggregate increase in the nonproduction labor share reflects skill biased technical change. However, our findings suggest that it will be difficult to find observable indicators of this skill biased technical change that account for the observed strong connection between the long run structural changes and cyclical dynamics. At least it will be difficult to find observable indicators from measures of investment in new capital goods (even very detailed indicators of the type of capital investment that the plant is undertaking) or R&D. The reason for this is that the latter indicators are procyclical and we find that the permanent changes in the nonproduction labor share are countercyclical.

An alternative but related interpretation of our results is that we have not fully captured the complex dynamic interactions between plant-level changes in the nonproduction labor share and plant-level changes in observable indicators like physical capital and R&D. For example, results from Dunne and Troske (1995) indicate that the probability of adopting advanced information technologies is increasing in the initial nonproduction labor share of the plant. We have not investigated the idea that plants may find it optimal to change their employment structure prior to changing their technology.⁵⁷ This latter possibility suggests that we may need a much richer characterization of the dynamic interactions of the various components of what constitutes the organizational structure of the plant.

⁵⁷ Note, however, that our long difference specifications relating long run changes in the nonproduction labor share to long run changes in output growth, physical capital growth and R&D capital growth do not depend on such precise timing considerations and these results still leave most of the aggregate and cross-sectional variation in the nonproduction labor share unexplained.

Appendix A: Definition of Measures

This appendix provides precise definitions for various measures used in the paper.

To begin with we define our job creation and destruction measures.

Job Creation by Worker Type

$$C_t^k = \frac{\sum_{i \in L^+} \Delta_{it}^k}{\sum_{i \in A} X_{it}^k}$$

Where:

Δ_{it}^k = Change in employment of worker type k at plant i in year t:
($L_{it}^k - L_{it-1}^k$)

$$X_{it}^k = 1/2(L_{it}^k + L_{it-1}^k)$$

L^+ = the set of all plants with $L_{it}^k \geq 0$.

A = All plants

Job Destruction by Worker Type

$$D_t^k = \frac{\sum_{i \in L^-} |\Delta_{it}^k|}{\sum_{i \in A} X_{it}^k}$$

Where:

L^- = the set of all plants with $L_{it}^k < 0$.

Job creation and destruction rates for total employment are calculated by letting L_{it}^k be the total employment in plant i in year t.

Given these measure of job creation and destruction, between plant job reallocation is simply the sum of total job creation and destruction:

$$\begin{aligned} BPR_t &= C_t + D_t \\ &= \frac{\sum_{i \in A} |\Delta_{it}|}{\sum_{i \in A} X_{it}} \end{aligned}$$

Total job reallocation (both within and between plant) is given by:

$$\begin{aligned}
 TR_t &= \frac{X_t^s}{X_t} (C_t^s + D_t^s) + \frac{X_t^u}{X_t} (C_t^u + D_t^u) \\
 &= \frac{\sum_{i \in A} (|\Delta L_{it}^u| + |\Delta L_{it}^s|)}{\sum_{i \in A} X_{it}}
 \end{aligned}$$

Finally, within plant reallocation is given by

$$\begin{aligned}
 WPR_t &= TR_t - BPR_t \\
 &= \frac{\sum_{i \in A} (|\Delta L_{it}^u| + |\Delta L_{it}^s| - |\Delta L_{it}|)}{\sum_{i \in A} X_{it}}
 \end{aligned}$$

Persistence Rate for Nonproduction Labor Share Change

Define:

$$\theta_{it}^m(k) = \max \left(\min \left(\frac{\sum_{j=0}^k \Delta M_{it+j}}{\Delta M_{it}}, \theta_{it}^m(k-1), 1 \right), 0 \right) \text{ for } k \geq 1$$

where

$$\begin{aligned}
 M_{it} &= L_{it}^s / (L_{it}^s + L_{it}^u) \\
 \theta_{it}^m(0) &= 1.
 \end{aligned}$$

$\theta_{it}^m(k)$ is the k-year persistence rate for plant i in year t. Then the k year persistent component of employment weighted, within plant nonproduction labor share change is given by:

$$\Delta M_t(k) = \sum_{i \in A} (L_{it-1}/L_{t-1}) \theta_{it}^m(k) \Delta M_{it}.$$

Persistence Rate for Job Creation

Define:

$$\theta_{it}^c(k) = \max \left(\min \left(\frac{\sum_{j=0}^k \Delta L_{it+k}}{\Delta L_{it}}, \theta_{it}^c(k-1), 1 \right), 0 \right) \text{ for } k \geq 0$$

where again $\theta_{it}^c(0)=1$ and calculated for plants with $\Delta L_{it} > 0$. Then the persistent component of aggregate job creation is given by:

$$C_i(k) = \frac{\sum_{i \in L^+} \theta_{it}^c(k) \Delta L_{it}}{\sum_{i \in A} X_{it}}$$

Rates and components for job creation by worker type, for total job destruction and for job destruction by worker type are defined in an analogous fashion.

Capital Stock Measures

Our capital stock measure is:

$$ks_{it}^r = (1 - \delta_{jt}^r) * ks_{it-1}^r + I_{it}^r$$

where:

ks_{it}^r = total stock of type r capital in plant i in period t
 δ_{jt}^r = depreciation rate for type r capital in industry j in period t (includes retirements).

I_{it}^r = Investment in type r capital in plant i in period t.

We construct capital stocks for t=1972 to 1988. For 1972, ks_{it}^r is computed as:

$$ks_{i,72}^r = bks_{i,72}^r * \frac{ks_{j,72}^r}{bks_{j,72}^r}$$

where:

$bks_{i,72}^r$ = the book value of type r capital in plant i in 1972

$ks_{j,72}^r$ = the real value of type r capital in industry j in 1972

$bks_{j,72}^r$ = the book value of type r capital in industry j in 1972.

The depreciation rate in period t along with the real and book value of capital for industry j in 1972 come from BEA. j refers to a two-digit industry.

R&D Stock Measure

Our R&D stock measure is:

$$rdk_{it} = (1-\delta)*rdk_{it-1} + rd_{it}$$

where:

rdk_{it} = the r&d stock of firm i in period t

$*$ = the depreciation rate for r&d

rd_{it} = the investment in r&d by firm i in period t .

We set $*$ = .15 (Adams and Jaffe, 1994). We compute r&d stocks for $t=1972$ to 1988. In 1972 we set $rdk_{it}=rd_{it}$.

Technology Adoption Measures

Our plant-level measure of technology adoption come from the 1988 Survey of Manufacturing Technology (SMT). The 1988 SMT surveyed managers of over 10,000 plants about their use of 17 different advanced technologies.⁵⁸ We merge these data with our sample of continuing plants to create a sample of 1820 plants. We then use the information from the 1988 SMT to construct our two measures of technology adoption: information, which equals 1 if a plant adopts any of one of eight information technologies, and production which equals 1 if a plant adopts any one of six production technologies. The eight information technologies are: computer aided design (CAD), CAD controlled machines, digital CAD, technical data network, factory network, intercompany network, programmable controllers, computers on the factory floor. The eight production technologies are: flexible manufacturing system/cell, materials working lasers, pick/place robots, other robots automatic storage/retrieval system, automatic guided vehicle systems, automated sensor based inspection and/or testing equipment used on incoming or in process materials, automated sensor based inspection and/or testing equipment used on final product.

Changes in Output, Ownership and Industry:

⁵⁸ See Dunne (1994) for a more complete description of the 1988 SMT.

Change in output is simply the log of the real value of shipments of plant i in year t minus the log of real value of shipments of plant i in year $t-1$. We deflate the reported value of shipments from the plant using the four-digit industry shipments deflator in the NBER productivity database (Bartelsman and Gray, 1995). In principle, the appropriate measure to use would be changes in value-added, or at least changes in shipments minus changes in inventories. However, given measurement error problems in the inventories and materials variables in the LRD and in the materials deflator, we feel that changes in shipments is the better variable. Similar arguments are made by Berman, Bound and Griliches (1994) and Bernard and Jensen (forthcoming) to justify their use of this variable.

The change in ownership variable equals one if the firm identification variable in the LRD changes between year t and year $t-1$. The change in industry variable equals one if a plant's four-digit SIC code changes between year t and $t-1$.

References

- Adams, James D. and Adam Jaffe, "The Span of the Effect of R&D in the Firm and Industry," Center for Economic Studies Discussion Paper no. 94-7. U.S. Bureau of the Census, Washington D.C., 1994.
- Altonji, Joseph G. and James R. Spletzer, "Worker Characteristics, Job Characteristics and the Receipt of On-the-Job Training," *Industrial and Labor Relations Review*, 45 (1991): 58-79.
- Amirault, Thomas, "Training to Qualify for Jobs and Improve Skills, 1991," *Monthly Labor Review*, 115 (1995): 31-36.
- Andolfatto, David and Glenn M. MacDonald, "Endogenous Technological Change, Growth, and Aggregate Fluctuations," mimeo, University of Rochester (1993).
- Bartel, Ann P. and Frank R. Lichtenberg, "The Comparative Advantage of Educated Workers in Implementing New Technology," *Review of Economics and Statistics*, 69 (1987): 1-11.
- Bartelsman, Eric J., Ricardo J. Caballero and Richard K. Lyons, "Customer and Supplier Driven Externalities," *American Economic Review*, 34 (1994): 1075-84.
- Bartelsman, Eric J. and Wayne Gray, "The NBER Manufacturing Productivity Database," mimeo, Board of Governors of the Federal Reserve System, (1995).
- Berman, Eli, John Bound, and Zvi Griliches, "Changes in the Demand for Skilled Labor Within U.S. Manufacturing Industries: Evidence from the Annual Survey of Manufacturing," *Quarterly Journal of Economics*, 109 (1994): 367-398.
- Bernard, Andrew B. and J. Bradford Jensen, "Exporters, Skill Upgrading, and the Wage Gap," *Journal of International Economics* (forthcoming).
- Berndt, Ernst, Catherine Morrison, and Larry Rosenblum, "High-Tech Capital Formation and Labor Composition in U.S. Manufacturing Industries: An Exploratory Analysis," Working Paper no. 4010, Cambridge, MA: NBER, 1992.
- Brown, Randall S. and Lauritis R. Christensen, "Estimating Elasticities of Substitution in a Model of Partial Static Equilibrium: An Application to U.S. Agriculture, 1947 to

- 1974," in Ernst R. Berndt and Barry C. Fields (eds.) *Modeling and Measuring Natural Resource Substitution*, Cambridge, MA: MIT Press, 1981.
- Caballero, Ricardo and Mohamad Hammour, "The Cleansing Effects of Recessions," *American Economic Review*, 84 (1994): 1356-1368.
- Caballero, Ricardo, Eduardo M. R. A. Engel, and John C. Haltiwanger, "Aggregate Employment Dynamics: Building from Microeconomic Evidence," Working Paper no. 5042: Cambridge, MA: NBER, 1995a.
- Caballero, Ricardo, Eduardo M. R. A. Engel, and John C. Haltiwanger, "Plant Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, no. 2 (1995b).
- Campbell, Jeffrey R. "Entry, Exit, Technology, and Business Cycles," Rochester Center for Economic Research Working Paper no. 407. University of Rochester, 1995.
- Chari, V.V. and Hugo Hopenhayn, "Vintage Human Capital, Growth, and the Diffusion of New Technology" *Journal of Political Economy*, 99 (1991): 1142-65.
- Cooley, T., J. Greenwood, and M. Yorukglu, "The Replacement Problem," mimeo, University of Rochester, Rochester, NY, (1994).
- Cooper, Russell and John Haltiwanger, "The Macroeconomic Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, 83 (1993): 360-82.
- Cooper, Russell, John Haltiwanger, and Laura Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," Working Paper no. 5260: Cambridge, MA: NBER, 1995.
- Davis, Steven J. and John Haltiwanger, "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," *NBER Macroeconomics Annual* no. 5 (1990): 123-68.
- Davis, Steve J. and John Haltiwanger, "Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986," *Brookings Papers on Economic Activity: Microeconomics*, (1991): 115-200.

- Davis, Steven J. and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107 (1992): 819-63.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh, *Job Creation and Destruction*, Cambridge, MA: MIT Press, 1996.
- Doms, Mark and Timothy Dunne, "Capital Adjustment Patterns in Manufacturing Plants," mimeo, Center for Economic Studies, U.S. Bureau of the Census, Washington D.C., 1994.
- Doms, Mark, Timothy Dunne, and Kenneth Troske, "Workers, Wages, and Technology," mimeo, Center for Economic Studies, U.S. Bureau of the Census, Washington D.C., 1995.
- Dunne, Timothy, "Patterns of Technology Usage in U.S. Manufacturing Plants," *The Rand Journal of Economics*, 25, (1994): 488-99.
- Dunne, Timothy and James Schmitz Jr., "Wages, Employment Structure and Employer-Size Wage Premia: Their Relationship to Advanced-Technology Usage at U.S. Manufacturing Establishments," *Economica* 62 (1995): 89-107.
- Dunne, Timothy and Kenneth Troske, "Human Capital, Research and Development Expenditures and the Adoption of New Technologies," mimeo, Center for Economic Studies, U.S. Bureau of the Census, Washington D.C., 1995.
- Dunne, Timothy, Mark Roberts, and Larry Samuelson, "Plant Turnover and Gross Employment Flows in the U.S. Manufacturing Sector," *Journal of Labor Economics*, 7 (1989): 48-71.
- Goldin, Claudia and Lawrence F. Katz, "The Origins of Capital-Skill Complementarity," mimeo, Harvard University, 1996.
- Griliches, Zvi and Jerry A. Hausman, "Errors in Variables in Panel Data," *Journal of Econometrics*, 31 (1986): 93-118.
- Hall, Robert E. "Labor Demand, Labor Supply, and Employment Volatility," *NBER Macroeconomics Annual* no. 6 (1991): 17-47.
- Hamermesh, Daniel S., *Labor Demand*, Princeton, NJ: Princeton University Press, 1993.
- Ichniowski, Casey and Kathryn Shaw, "Old Dogs and New Tricks: Determinants of the Adoption of Productivity-Enhancing Work

- Practices," *Brookings Papers on Economic Activity: Microeconomics*, (1995): 1-66.
- Jovanovic, Boyan and Glenn M. MacDonald, "Competitive Diffusion," *Journal of Political Economy*, 98 (1994): 24-52.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce, "Wage Inequality and the Rise in the Return to Skill," *Journal of Political Economy*, 101 (1993): 410-42.
- Juhn, Chinhui, Kevin M. Murphy and Robert H. Topel, "Why Has the Natural Rate of Unemployment Increased Over Time?" *Brookings Papers on Economic Activity* no. 2, (1991): 75-142.
- Katz, Lawrence F. and Kevin Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107 (1992): 1-34.
- Kremer, Michael. "The O-Ring Theory of Economic Development," *The Quarterly Journal of Economics*, 108 (1993): 551-575.
- Lambson, Val E., "Industry Evolution with Sunk Costs and Uncertain Market Conditions," *International Journal of Industrial Organization*, 9 (1991):171-96.
- Lynch, Lisa M. and Sandra E. Black, "Beyond the Incidence of Training: Evidence from a National Employers Survey," Working Paper no. 5231. Cambridge, MA: NBER, 1995.
- Mortensen, Dale T. and Christopher A. Pissarides, "Job Creation and Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61 (1994): 397-416.
- Shea, John, "Do Supply Curves Slope Up?" *Quarterly Journal of Economics*, 108 (1993): 1-32.
- Solow, Robert M., "Investment and Technical Progress," in Kenneth Arrow, Samuel Karlin, and Paul Suppes (eds.), *Mathematical Methods in the Social Sciences*, 1959, Stanford, CA: Stanford University Press, 1960.
- Troske, Kenneth R. "Evidence on the Employer Size-Wage Premium from Worker-Establishment Matched Data," Center for Economic Studies Discussion Paper no. 94-10. U.S. Bureau of the Census, Washington, D.C., 1994.

Troske, Kenneth R. "The Worker-Establishment Characteristics Data," Center for Economic Studies Discussion Paper no. 95-10. U.S. Bureau of the Census, Washington, D.C., 1995.

U.S. Census Bureau, 1987 Census of Manufactures, General Summary, MC87-S-1, Washington, D.C.: Government Printing Office, 1991.